

Multimodal non-invasive hemodynamic monitoring

by

Barry Belmont

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Biomedical Engineering)
in the University of Michigan
2016

Doctoral Committee:

Professor Albert J. Shih, Chair
Professor Joseph Bull
Professor Chinedum Okwudire
Professor Kevin R. Ward
Professor William F. Weitzel
Professor Zhen Xu

©Barry Belmont

2016

ACKNOWLEDGMENTS

We never really go it alone. During my time at the University of Michigan I have had the absolute pleasure of working with some of the best colleagues and collaborators a person could ever ask for. Each in their own way has contributed to this dissertation, making this work possible, and helping make me who I am today and what I will become tomorrow.

I must first and foremost thank Hakam Tiba, without whom perhaps none of this work would have come to fruition. His patience and perseverance when working with a rather obtuse and sometimes obstinate engineer were nothing short of saintly. Thanks must also go to the rest of the core team of emergency medicine collaborators I worked with regularly—Jerry Draucker, Amanda Pennington, Brandon Cummings—whose early mornings and long days were more than I could ever ask for.

The consistent help of many emergency medicine physicians, most especially Nik Theyyuni, Chris Fung, Rob Huang, and Ross Kessler, both in explaining and in practicing their craft was crucial to what I have reported here.¹ These are people whose job it is to save lives and with what little free time they had to spare helped examine hundreds of patients because they thought this work could help make a difference. This dissertation stands as a testament to such individuals in general and to these individuals in particular.

In my lab, there are far too many people to whom I owe far too much. From the old guard—Roland Chen, Rob Dodde, Grant Kruger, and Bruce Tai—to the next generation—Annie Li and Robert Chisena—working at Wu Manufacturing Research Center has been one of the most rewarding experiences of my life. The outsized work ethic, output, and helpfulness of Lei Chen, Yuan Jin, Ali Kuzu, Hui Li, Weisi Li, Jeff Plott, and Yihao Zheng cannot be overstated. This is the hardest working, smartest, and kindest bunch of people you will ever find.

During my time there I was fortunate enough to call several students mentees. Five I would like to specially acknowledge are Adetunji “TJ” Dahunsi, Luc Hildebrand, Il Taek Kwon, Carline Mayo, and Tong Xie who put up with far more of my nonsense than anyone should be expected to. In frustration and exhilaration, getting to see them grow as individuals and engineers was a profound honor.

Outside of the research grind I was lucky enough to find a few folks that shared my passion for teaching. I would like to particularly thank a few members of the University of Michigan’s student chapter of the American Society for Engineering Education—Aaron

¹It should perhaps be noted here that any and all errors in the physiology to follow are wholly my own. These folks really know what they’re talking about. They’re real doctors. I, at the time of this writing, am not quite there.

Bevill, Andrew McAllister, Cameron Louttit, Ibrahim Mohedas, John Pitre, and Kathleen Ropella—who somehow put up with many of my harebrained schemes, while also promoting engineering education.

In the same vein, I would also like to thank each and every one of my students for helping me learn how to teach (so far as I am able): Bora Agabigum, Tom Angers, Oliver Bloom, Nick Caccamo, Qin Chen, Bailey Cockerill, Braden Deweese, Brian Downey, Patrick Dunlop-Karagianis, David Feddema, Matthew Federbush, Rachel Ferguson, Kaity Forde, Zachary French, Justin Gasparovich, Stephanie Goretski, Melanie Hamilton, Christina Healy, Umakanthan Kavin, Doga Kumusoglu, Melissa Lemke, Harry Liu, Tony Lu, Yi Luan, Andrew Lynch, Emily Magurno, Bethany Meyer, Ian Moore, Chad Newton, Andrew Nikolai, Anastasia Ostrowski, Haresh Patel, Vivek Pejaver, Maggie Pistella, Julia Prisby, Kevin Renehan, Danika Rodrigues, Katelyn Rowley, Nabeel Salka, Natalie Setterbeg, Paul Snyder, Lisa Solomon, Samantha Spierling, Shriram Srinivasa-Narasimhan, Chris Stephan, Jordan Sykes, Abishek Thiagaraj, Ryan Thomas, Kyle Thompson, Walker Thompson, Drew Vazquez, Eric Westfall, and Kyle Wolok. I hope to have had at least a fraction of the impact on each of them as they all did on me.

I am also thankful to many people defying typical classification in these sorts of things: Karen Schroeder for a friendship from the beginning; Mark Stratton for a belief in me; Alan Vollmer for an ever ready helping hand in clinic; and each and every one of the nurses and technicians for pointing me in the right direction and telling me when I had got something wrong.

A man is nothing without his family. I am no exception. My mom and dad, Mary and Bob Belmont, built me up with love and discipline. My sister, Kelli Belmont, has been with me through it all, from the tragic to the absurd. May this work reflect each of these attributes.

Though much of the work that follows is on the function of the heart, Julia Kline taught me more on the subject than I can report here. Physiologically mine may be no different for having met her, existentially it will never be the same.

I must also express my gratitude for each of the members making up my committee. Joe Bull and Zhen Xu have been wonderful academic and scholarly advisors, helping when and where they could to make my research better than it might otherwise have been.

A last minute addition to my committee, Chinedum Okwudire has shown me what it is to be a dedicated young researcher. Driven, smart, charming. We may all take solace for manufacturing's bright future as this man becomes one of its leading researchers.

Kevin Ward and his Center for Integrative Research in Critical Care (MCIRCC) have changed not only how I see my science (as a means to laudable ends), my pursuits, and my

approaches, but has shown me there is no surer way to change something than determination and collaboration. If it can be said that I have bore fruit with this work, it is in large part due to the fertile ground MCIRCC has provided me.

Pick a positive word to say, pick another, and continue until you run out of words. Each one them applies to Rick Weitzel in spades. Thoughtful, kind, observant, clever, Dr. Weitzel is everything you want out of a physician, a researcher, and a person. That he happens to be one of the best for all three should escape no one's attention.

My chair, my professor, my boss, my mentor, Albert Shih is the best advisor I could have ever asked for. He took a chance on a kid from Nevada when no one else would. He has given me more opportunities than I ever deserved and if what I have reported here is at all good, it is from his guidance.

Finally, I would like thank each and every patient and their families for trusting me to make something of the data they shared. I can only hope this dissertation goes some small way towards repaying them.

TABLE OF CONTENTS

Acknowledgments	ii
List of Figures	vii
List of Tables	xviii
Abstract	xx
Chapter	
1 An introduction to hemodynamics	1
1.1 The heart	4
1.2 The vasculature	14
1.3 Blood flow	17
1.4 Renal function	23
1.5 Volume status	28
1.6 Monitoring	29
1.6.1 The physical exam	30
1.6.2 Static measures of intravascular volume	30
1.6.3 Pressures in general	31
1.6.4 Volumes in general	36
1.6.5 Dynamic measures of intravascular volume	39
1.6.6 Variations in general	40
1.7 A summary	48
2 Ultrasound monitoring	50
2.1 Ultrasound imaging in general	50
2.2 Theoretical foundations for ultrasound imaging	56
2.2.1 The wave equation and its solutions	56
2.2.2 Acoustic impedance and wave reflections	60
2.2.3 Attenuation and enveloping	64
2.2.4 Envelope detection and image reconstruction	66
2.3 Basics of image registration	69
2.3.1 Feature detection	70
2.3.2 Feature matching	71
2.3.3 Image transformation	74
2.3.4 Iterative Kanade-Lucas-Tomasi feature tracking algorithm	83
2.4 Implementation	90

2.4.1	Example: Arteriovenous fistula mechanical parameters and the development of an open-source ultrasound toolbox	90
2.4.2	Example: Continuous inferior vena cava tracking	109
2.4.3	Conclusions	117
3	Bioimpedance monitoring	118
3.1	Theoretical basis of bioimpedance	118
3.1.1	Permittivity and polarization	119
3.1.2	Complex permittivity	123
3.1.3	Relaxation and dispersion	125
3.1.4	Biological materials and their electrical equivalents	126
3.2	Practical application of bioimpedance	132
3.2.1	Clinical use	132
3.2.2	Proposed use	136
3.3	Impedimetric intravascular volume evaluation	141
3.3.1	Experimental evaluation	144
3.3.2	I2VE of control and hospitalized subjects	153
3.3.3	I2VE during hemodialysis	155
3.3.4	I2VE in mechanically ventilated subjects	159
3.3.5	A few limitations of I2VE	162
3.3.6	Other observations	163
3.3.7	Conclusions and the goals of future work	167
4	Wearable bioimpedance monitoring	168
4.1	The rise of wearable monitoring	168
4.1.1	Biomedical signals in context	168
4.1.2	The potential role of bioimpedance	171
4.2	A wearable bioimpedance monitor	173
4.2.1	The design	175
4.2.2	The validation	193
4.2.3	The results	195
4.2.4	Conclusions	199
5	The shape of things to come	200
	Bibliography	205

LIST OF FIGURES

1.1	The circuits of the circulatory system. In human beings the pulmonary circuit begins at the right ventricle where blood is ejected to the lungs and ends at the left atrium where oxygenated blood enters the heart. The system circuits begins with blood being ejected from the left ventricle and with it returning through the right atrium. Image is from http://msjensen.cehd.umn.edu/Webanatomy/imagetdatabase/Cardiovascular/circulation.tif and is used for educational purposes.	5
1.2	Basic anatomy of the heart demonstrating the relationships between the atria, the ventricles, and the vessels what feed into them. The flow of blood within the heart is indicated with arrows. Modified from "Diagram of the human heart (cropped)" licensed under CC-BY-SA-3.0 via Wikimedia Commons.	6
1.3	Pressure and volume relationships during systole and diastole. Modified from "Wiggers Diagram" by DanielChangMD revised original work of DestinyQx; Redrawn as SVG by xavax - Wikimedia Commons, File:Wiggers Diagram.png.. Licensed under CC BY-SA 2.5 via Wikimedia Commons	9
1.4	Schematic representation of the main factors affecting cardiac output. Cardiac output, typically represented as the product of stroke volume and heart rate, is a measure of the volumetric flow rate of blood pumped out of the ventricles to the pulmonary and systemic circuits. It is affected and affects many hemodynamic parameters and its relationship to the larger fluid volume management capabilities of the body must be understood in context.	10
1.5	The cardiac function curve as predict by the Frank-Starling law and its coupling with venous return. The intersection of the two curves represents the operating point of the heart. Changes to either the cardiac output or blood volume curve affect how much blood can be circulating and in what way the heart circulates that blood.	13
1.6	The relationship between the (a) diameter, (b) cross-sectional area, (c) average blood pressure, and (d) velocity of blood with respect to their vessel type in the cardiovascular system.	15
1.7	A schematic representation of the kidney. Blood flows from the renal artery into the segmental arteries to the interlobar arteries to the acruate arteries to the interlobular arteries until finally reaching the capillary beds of individual nephrons. Once there filtration can begin with excess fluid removed, traveling from the ureter into the rest of the urinary system.	24

1.8	A simplified representation of renal function. Blood traveling through the afferent arteriole enters the glomerular capillaries in Bowman’s capsule, where filtration occurs. Blood passing from the efferent arteriole reabsorbs and secretes fluids and materials in the peritubular capillaries, draining finally into the renal vein. All filtered fluid not reabsorbed is excreted.	25
1.9	Central venous pressured via insertion of a central venous catheter, also known as a ”central line.” The tip of the pressure transducer is fed through the vasculature into the superior vena cava until it is placed just outside the right atrium. A peripherally inserted catheter, or a ”PICC line,” can also be used to obtain central venous pressure measurements, using the peripheral vasculature to guide the pressure transducer tip to the same point in the vena cava.	33
1.10	Fed through a peripheral vein, in this case the femoral, the balloon tip of the Swan-Ganz catheter is guided through the right atrium into the right ventricle up through the pulmonary artery. When the balloon is inflated, blood flow is obstructed and the resulting pressure difference on either side, the pulmonary artery occlusion pressure, results from the pressure of the left atrium. The end-diastolic pressure in the atrium is treated as an equivalent of the pressure (and volume) in the left ventricle, and thus cardiac preload.	34
1.11	Right ventricular end-diastolic volume represent the amount of blood left in the right heart at the end of diastole. The volume of the right ventricle at the end of diastole is difficult to measure with bedside techniques and is thus usually back-calculated from thermodilution based methods for determining ejection fraction.	36
1.12	The inferior vena cava (IVC) as measured through subcostal echocardiography. The heart, right atrium (RA), liver, hepatic vein (HV), and common bile duct (CBD) all provide landmarks by which to make a consistent IVC diameter measurement.	38
1.13	A typical arterial waveform collected using photoplethysmography. The respiratory induced variation in the waveform is used to determine a subjects volume status.	42
1.14	Respiratory induced variations in inferior vena cava diameter as imaged through subcostal echocardiography. (Top) B-mode ultrasound images of the inferior vena cava during inspiration (left), where the diameter is at a minimum, IVC_{min} and expiration (right), where the diameter is at a maximum, IVC_{max} , in a spontaneously breathing individual. Taking an A-line distal of the hepatic vein produces an M-mode image where the respiratory (and even cardiac) variations can be seen.	44
1.15	The procedure for performing a passive leg raise. First a subject is placed in a semi-recumbent position (at 45). After a short period of time, the patient is laid supine and their legs are raised 45 of the horizon. In this way blood from the legs is pushed to the heart, performing a non-invasive fluid challenge. . . .	46

2.1	A schematic representation of a typical clinical ultrasound imaging system. The signal for a sound wave to be injected into an object through a transducer array is produced by a central controller, modulated, converted to an analog signal (through a digital-to-analog converter (DAC)), and amplified through a transmit beam former. The echoes from the input waves are received by the transducer, amplified, and converted into a digital signal (through an analog-to-digital converter (ADC)) via a receiver beamformer. That digital signal is then used to estimate the envelope of the signal for further image processing and display.	52
2.2	Brightness mode (B-mode) ultrasound example of a post-arterial anastomotic segment of an arteriovenous fistula connected to a cephalic vein. The areas of darkness are caused by an absence of specular reflection typical of blood. The single straight line down the center of the vessel, known as an A-line, was placed by a researcher there to elucidate how A-mode and M-mode images are found (see Figures 2.3 and 2.4).	53
2.3	An example of amplitude mode (A-mode) ultrasound. The amplitude of the signals reflected from the medium are plotted against the time at which the signal was received. These amplitudes are taken from the A-line of Figure 2.2. Note that because of ultrasounds high frequency content (in this case the transducer used had was 7.5 MHz), the time over which these signals are collected is often less than a millisecond (and in this case is approximately 35 μ s).	54
2.4	A motion mode (M-mode) ultrasound image taken from the A-line seen in Figure 2.2. Here the slight pulsatility of the vessel can be observed in a manner that would not obvious from a single B-mode image (as seen in Figure 2.2) or A-mode plot (as seen in Figure 2.3). Another example (in a more dynamically varying vessel) can be seen in Figure 1.14.	54
2.5	(Top) Doppler mode image of the arteriovenous fistula seen in Figure 2.2. A color gradient (not shown) specifies the velocity measured by the dephased acoustic waveform. (Bottom) A velocity profile can be reconstructed from the doppler image and plotted against time. Clinicians will often integrate a velocity profile found from a doppler mode ultrasound to calculate the volumetric flow.	55
2.6	A small volume of matter (left). When a pressure is applied to the y-z surface, the x dimension is increased by an amount Δx (center), the left most surface of the element is moved by a distance u and the right most surface is moved by a distance $u + \Delta u$ (right).	57
2.7	A small volume of matter (top). When a pressure is applied to the y-z surface, the x dimension is increased by an amount Δx (middle), the left most surface of the element is moved by a distance u and the right most surface is moved by a distance $u + \Delta u$ (bottom).	62
2.8	Gaussian pulse modulation of a signal. A carrier frequency signal (top) is modulated by a Gaussian envelope (middle) to produce a pulse resistant to depth dependent blurring.	65

2.9	An example of a Gaussian pulse signal subjected to a Hilbert transform. Note that what appears to be an envelope around the signal was found by taking the absolute value of the Hilbert transform, thus demonstrating its robust ability to estimate the envelope of a signal.	67
2.10	Ultrasound waveform envelope transformed into pixel intensity values. This data represents the first 6 mm (approximately $8 \mu s$) of the A-mode ultrasound waveform shown in Figure 2.3. Such a process is done for all data received by the transducer to produce a sequence of two-dimensional images. What such pixel values look like in practice may be seen in the top quarter of the M-mode image of Figure 2.4.	68
2.11	Correlation-based feature matching. (a) A subset of an image, called the kernel, is selected (in this case a portion of the post-arterial anastomotic segment from the fistula ultrasound originally shown in Figure 2.2) and its cross-correlation to every position within the search window (in this case, the entire image) can be calculated. (b) A map of the correlation coefficients at every position in the window is shown, and indeed the area of maximal correlation corresponds to the original subset image. (Red represents high values, blue represents low values.)	72
2.12	Image pyramid of an ultrasound frame looking at a subject's inferior vena cava created by subsampling segmentation.	86
2.13	An basic example of the iterative process used to track a feature across a pyramidally segmented image. Here an original image, at a level L_0 , is downsampled twice. Once at its highest level, L_2 , initial tracking begins. The results from this tracking are then fed forward to levels L_1 and L_2 , with the results of each, d^{L_m} , adding to final displacement vector, \mathbf{d}	88
2.14	Determination of feature selection appropriateness. (a) A map of the eigenvalue pairs values (with red being the highest and blue being the lowest). (b) The top 500 points of maximum value from (a) plotted atop the image from which they were taken. Notice the lack of features seen in the blackness of the blood.	89
2.15	Anatomy of the upper limb. Several blood vessels vessels relevant to arteriovenous fistulas (brachial, radial, and ulnar arteries; axillary, brachial, basilic, and cephalic veins) are highlight. (Modified from the <i>Atlas of Dialysis Vascular Access</i> for educational purposes.)	92
2.16	A few major types of arteriovenous fistulas. These fistulas are formed by the pairing for the the (a) cephalic vein and radial artery, (b) perforating branch and proximal radial artery, (c) median antecubital vein and radial artery, and (d) cephalic vein and brachial artery. (Modified from the <i>Atlas of Dialysis Vascular Access</i> for educational purposes.)	93
2.17	Two types of peripheral graft-based vascular accesses: (a) forearm loop and (b) upper arm. In both cases a synthetic tube redirects arterial blood to the venous side to modify the size, shape, and physical properties of the veins. (Modified from the <i>Atlas of Dialysis Vascular Access</i> for educational purposes.)	94

2.18	Main GUI of the software toolbox developed to aid in fistula maturation research. Its main features from this window include (a) viewing the color histogram of the image from which (b) a user can manually set brightness and contrast parameters, (c) a zoom-to-fit function to get the most use out of the space, (d) overlaying a grid on top of the image, (e) the ability to change the colormap, (f) a save function for the marked-up ultrasound image in the window as either (g) a single frame or the entire video, (h) the ability to edit or export data made by the user during a session, (i) importing a new image to work with, (j) displaying the current frame number and the total number of frames, (k) applying either a median or Kalman filter across the frames to minimize certain types of noise in the images, (l) setting up the image properly (such as inputting or calibrating the distance between pixels, determining the number of pixels to be analyzed during strain measurements, establishing the bounds for accumulated strain, inputting blood pressure values, specifying the frame rate, or asking for help), (m) analyzing the pulsatility, strain (either total image or vessel wall specific), or blood-wall shear rate of the image, (n) manipulating the image for analysis (such as cropping (either manually or automatically), mapping pixel intensities for a given perimeter, measuring point-to-point distances, or removing frames), (o) specifying regions of interest, (p) displaying mean and standard deviations of certain measurements within regions of interest, (q) displaying the pixel intensity values at the cursor, and (r) scrolling through each of the video frame by frame.	95
2.19	The (a) diameter and (b) distensibility of the vessel measured in Figure 2.2 over time. On the left-hand y-axes the absolute measures of diameter and distensibility can be seen, while on the right hand y-axes the relative measures of diameter (from minimum) and distensibility (across range) are shown. . . .	98
2.20	Strain map development of a radiocephalic fistula over time. The magnitude of strain is presented here (a) just after the initiation of a cardiac cycle, (b) halfway to peak strain, and (c) peak strain. The pixel intensity represents the magnitude of strain and local maximums can be found at both the arterial wall and one edge of the post-arterial anastomotic segment. (These images were produced by subsampling the original image by 20%.)	99
2.21	Blood velocity profile in a brachial artery at peak systole from both a decorrelation measurement and as an ideal Poiseuille flow.	100
2.22	The initial segmented vessel geometry (left), a smoothed and splined edge geometry (middle), and the final solid three-dimensional geometry (right). Only 1% of available edge pixels are displayed in the initial segmented image to make its properties clear. Similarly, every other image slice is shown for the smoothed and splined geometry. The final three-dimensional geometry here was used for some of the fluid simulations reported in Figures 2.23-2.25. . . .	102
2.23	Velocity streamlines across the simulated fistula.	104
2.24	Wall shear stress across the simulated fistula.	105
2.25	Wall shear rate across the simulated fistula.	105

2.26	Minimum diameter of the inflow artery, post arterial anastomotic segment (PAAS), and the outflow vein of measured during an initial evaluation (n = 24 patients). The minimum diameter is reported here as the parameter of interest, distensibility—the amount a vessel expands from a baseline—relies on it. While the distributions of the vessel sizes are fairly normal, a few outliers exist in the veins with a pair of measured veins very large and a single small vein. The patients to whom these outliers belong were removed from the follow-up examination evaluation (seen in Figure 2.28).	106
2.27	Distribution of (a) pulsatile differences and (b) distensibilities for the inflow artery, post arterial anastomotic segment, and outflow vein measured during an initial evaluation (n = 24 patients).	107
2.28	Clinical data relating the change in (a) vessel diameter and (b) distensibility measured during an initial evaluation (approximately one week post-operation) and a follow up examination (typically one to two months post-operation). Note the increased size of the vein and the increased pulsatility of the artery in the follow up evaluations as compared to the initial evaluations.	107
2.29	Compliance of the artery, post-arterial anastomotic segment, and vein during the initial evaluation and the follow up examination. (Data presented is of ten patients with complete records.)	109
2.30	Varieties of inferior vena cava images. A point tracking system must be robust enough to handle the substantial variety of ultrasound images that arise from patients.	112
2.31	Representative example of the inferior vena cava diameter tracked over time through ultrasound and the iterative KLT algorithm. Note that the raw signal (top) is unfiltered and represents the motion of the inferior vena cava precisely. Removing the respiratory signal produces a cardiac only signal (bottom) whose amplitude varies with respiration—increasing during inspiration and decreasing with expiration.	114
2.32	A histogram of (a) the maximum diameters of the IVCs as measured by the clinicians and (b) the caval index of the IVCs as measured by the clinicians. While a fairly normal distribution was found for maximum diameters (centered around 18 mm), the caval indices demonstrated a skewed distribution with a long tail tending toward 100% collapse.	114
2.33	Bland-Altman plot (mean of measurements vs. the difference in measurements) of the caval indices as measured by physicians and the algorithm. The three white points at high levels of collapse represent outliers in the data set.	115
3.1	Point charges demonstrating Coulomb’s law.	119
3.2	Simple capacitor representation of dielectric material.	123
3.3	Dielectric responses. Modified from ”Dielectric responses” by Archimerged found at Wikipedia. Licensed under Attribution via Commons - https:// commons.wikimedia.org /wiki/	125
3.4	Idealized behavior of dielectric dispersions for biomaterials showing α , β , and γ regions. Modified for educational purposes from <i>Bioimpedance and Bioelectricity Basics</i> by Grimnes and Martinsen, originally from Schwan.	127

3.5	The cells of a tissue and their electrical component equivalents. For an individual cell, R_m and C_m are the resistance and capacitance of the cellular membrane, R_i is the intracellular resistance, and R_e is the extracellular resistance. Summing up all the cells yields one of two popular equivalent circuits to work with, one (right top) the parallel combination of a resistor, R_1 with a series combination of a resistor, R_2 with a capacitor, C , and the other (right bottom) the parallel combination of a resistor, R_1 with a series combination of a resistor, R_2 with a constant phase element, CPE	129
3.6	Representative resistance-reactance plot for biological tissues. $\mathbf{Z}(\omega)$, $\theta(\omega)$, $R(\omega)$, and $X(\omega)$ are the impedance, phase angle, resistance, and reactance at an angular frequency, ω , R_0 is the resistance at low frequencies, R_∞ is the resistance at high frequencies, X_{peak} is the peak reactance and it occurs vertically over the depressed locus, D	130
3.7	Representative bioimpedance vector analysis graph. The impedance vector normalized to the patient height falls within a probabilistic region to determine the ratio of fluid to soft tissues a person has.	134
3.8	An example of electrical impedance tomography demonstrating (a) healthy and (b) unhealthy lungs. The pixel intensities here represent a range-scaled map of electrical impedance measured across sixteen electrodes along the subject's chest (represented by the perimeter of white dots in each image). In the case of the healthy lungs, air (and thus impedance) is evenly distributed between the front and back of both lungs. The unhealthy lungs, from a subject under mechanical ventilation undergoing tidal recruitment, show one lung and one portion of that lung preferentially letting in air. The images of this figured were generated using data supplied by Drager.	136
3.9	A representation of the parameters involved in tetrapolar impedance measurement. Four electrodes, A , B , C , D are placed on a volume, v , with potentials ϕ and ψ , current densities \mathbf{J}_ϕ and \mathbf{J}_ψ , and conductivities σ_ϕ and σ_ψ . A bolus of fluid travels from A to B causing a conductance change, a potential difference, and a measurable impedance variation across the electrodes.	139
3.10	A representative spontaneous breathing baseline. Here a range normalized bioimpedance waveform (solid line) demonstrates two distinct components: a higher frequency and often lower amplitude signal is superimposed upon a lower frequency corresponding to respiration (dashed line).	142
3.11	An example of a bioimpedance waveform collected during forced inspiration. Two normal breaths are preceded by four forced breaths (each limited to 20 cm-H ₂ O [2 kPa]). A distinct positive and magnified peak-to-peak amplitude difference can be seen compared to the normal inspiratory maneuvers preceding them.	143
3.12	An upper limb bioimpedance waveform during forced expiration. Two breaths outward (each limited to 10 cm-H ₂ O [1 kPa]) have a negative and magnified peak-to-peak amplitude with respect to the normal breathing baseline preceding them.	144

3.13	Electrode placement along the upper limb for the I2VE experiments reported here. Two current injecting electrodes reside near the shoulder and the wrist/forearm and two potential sensing electrodes were placed closer to each other along the bicep. Image from US 2015/0031966 A1 and is used for educational purposes.	149
3.14	A typical impedance waveform observed before and after a forced respiratory maneuver demonstrating each of the the I2VE parameters: the baseline peak to peak impedance change due to spontaneous breathing, Z_b ; a series of forced inspiration causing individual amplitude changes, $Z_{m,1-3}$; a potential change in the baseline amplitude after the maneuver, $Z_{b,a}$; and a possible baseline shift, Z_{bs} , increasing the impedance minimum, distinct from the amplitude change. The I2VE parameter, dZ , may be calculated from equation 3.48. Note that Z_b here is shown as a single wave for convenience. In reality the unperturbed breathing baseline is an average value of many waveforms.	152
3.15	The results of the volunteers and emergency room patients: (a) inferior vena cava diameter collapse vs. impedance change from baseline; (b) ROC curves of several single threshold tests; (c) area under the ROC curves for every collapse threshold observed (20-100%); and (d) area under the ROC curve as a function of the lower and upper bounds of the multiple threshold analysis.	154
3.16	Inferior vena cava diameter collapse vs. impedance change from baseline (a) before and (b) after dialysis. Both the before and after measurements show moderate correlation (R^2 values of 0.664 and 0.657) for their linear regression models of $dIVC = 35.9 \ln(dZ) + 16.2$ and $dIVC = 41.1 \ln(dZ) + 13.8$, respectively.	156
3.17	Example ROC curves of five thresholds of inferior vena cava collapse (10%, 30%, 50%, 70%, 90%) (a) before and (b) after dialysis. Areas under the curve in these examples ranged from 0.88 to 0.98 before dialysis and from 0.88 to 0.96 after, demonstrating good predictive power in both scenarios.	157
3.18	Area under the ROC curves for collapse thresholds from 10-100% in subjects both (a) before and (b) after dialysis. A peak of 0.979 was observed at 50% collapse before dialysis, while two potential peaks of 0.975 and 0.959 appear at 33% and 61% after.	158
3.19	Areas under the ROC curve as a function of the lower and upper collapse threshold (a) before and (b) after dialysis. Large sections of both graphs show AUCs greater than 0.99, with the extrema seeing the lowest values (still above 0.85 in all cases before dialysis and above 0.88 in all cases after).	158
3.20	The total results from patients receiving hemodialysis (both before and after treatment): (a) inferior vena cava diameter collapse vs. impedance change from baseline (showing moderate correlation ($R^2 = 0.69$) for its linear regression model of $dIVC = 38.08 \ln(dZ) + 15.48$); (b) ROC curves of several single threshold tests (with AUCs ranging from 0.87 to 0.95); (c) area under the ROC curves for collapse thresholds from 10-100% (with a peak value of 0.968 at 33% collapse); and (d) area under the ROC curve as a function of the lower and upper bounds of the multiple threshold analysis (with the region between the lower thresholds of 10-30% and upper thresholds of 30-100% yielding the best results (AUCs >0.95)).	160

3.21	The total results from patients under mechanical ventilation: (a) inferior vena cava diameter distension vs. impedance change from baseline (showing moderate correlation ($R^2 = 0.57$) for its linear regression model of $dIVC = 62.15(dZ) - 47.49$); (b) ROC curves of several single threshold tests (with AUCs ranging from 0.73 to 0.99); (c) area under the ROC curves for distension thresholds from 1-100% (with a dip to 0.67 at the 6% threshold); and (d) area under the ROC curve as a function of the lower and upper distension bounds with AUC increasing as a function of increasing of the upper limit threshold.	161
3.22	The arterial and venous components of impedance plethysmography. (a) The original signal and the respiratory signal, the latter of which represents venous blood movement. (b) The cardiac signal of impedance separated from the original and respiratory signals. Note the relative amplitude is approximately $\pm 10\%$ that of the total signal. (c) The first order derivative of the cardiac impedance signal over time, an approximate stroke volume over time measurement.	164
3.23	Representative arterial waveform found via impedance plethysmography by taking the first order derivative of the cardiac signal. Characteristic features seen in other arterial monitoring methods, such as systolic uptake, systolic decline, and the dicrotic notch are present here.	165
3.24	Impedance over time of a sleeping, snoring patient. The subject takes a series of large and rapid breaths before settling down towards a shallow breathing baseline, only to start the cycle anew.	166
3.25	Decreasing respiratory rate as measured from the second half of Figure 3.24. Being able to track the breath to breath accuracy with impedance could open up many potential avenues of research.	166
4.1	A block diagram representation of the wearable bioimpedance monitor developed. Electrodes are placed on an objected and connected via a 3.5 mm audio jack to the device. The device itself then measures the impedance, sends that data to a central processor, then transmits that data either wirelessly (Bluetooth) or through a USB connection. The USB connection also serves to program the device and charge the battery. Each of step and component of this process is discussed in this chapter. The actual printed circuit board based upon this diagram used in this work can be seen in Figure 4.2(b).	174
4.2	A (a) simplified representation of the printed circuit board and (b) the actual printed circuit board itself. In the top left corner is a button flanked, in the top right is another button and LED combination, the battery connector, and the sliding switch to turn the device on and off. Towards teh middle of the board is a reset button, microcontroller, and AD5933 and its analog front end. At the bottom from left to right is a 3.5 mm audio jack connection, a micro-USB connector, and the Bluetooth antenna. (Actual size if printed on 8.5" x 11" paper.)	176

4.3	Operation of the AD5933. A sinusoidal voltage signal at a designated frequency is generated via the combination of the master clock (MCLK), an oscillator, and a direct digital synthesizer (DDS). This discrete signal is then converted to analog and amplified, reaching the material as VOUT. Impedance in the material, $Z(\omega)$, is found by passing the current at VIN through an amplifier, filtering, converting to a digital signal, then taking the discrete Fourier transform at the digital signal processing (DSP) engine with information supplied by the DDS core. The real and imaginary components of the impedance are then passed to an inter-integrated circuit (I2C) communication interface and transmitted out through serial data (SDA) and serial clock (SCL) lines.	178
4.4	The reconfigurable frontend designed to work with the AD5933 shown in Figure 4.3. Such a front end allows a bipolar or tetrapolar electrode arrangement to be specified on the fly, enabling researchers to choose that which best fits their needs.	180
4.5	The processor (ATmega1280) of the wearable device and each of its relevant connections. Each of the main functions of the device has a presence here: communication with the AD5933 (PD0 and PD1), USB (PE0 and PE1), and Bluetooth (PH0 and PH1); memory storage (PH6 and PH7); and programming (such as PB1-3).	181
4.6	I2C voltage translator to make possible communication between the microcontroller powered at 3.3 V and the AD5933 powered at 5 V. Signals from the microcontroller, SDA and SCL, are stepped up to become SDA_5V and SCL_5V, to interact with the impedance analyzer.	184
4.7	USB connection and the battery charging circuit. Potential from the USB port charges the battery which in turn powers the device. Signals from D- and D+ (USBDM and USBDP) are sent to the USB to serial converter seen in Figure 4.9.	185
4.8	Battery regulator components of the circuit ensuring that consistent potentials of 3.3 and 5 V are available to those components of the device that require them. Both the linear regulator and the voltage converter are designed with the low quiescent currents of wearable electronics in mind.	186
4.9	USB to serial converter that enables wired communication—be it programming, data transmission, or data reception—between the device and a computer.	187
4.10	Bluetooth radio.	188
4.11	Accelerometer for actigraphy monitoring, motion activation, and compensation.	189
4.12	Local storage of measured data before transmission. Over two full days of complete impedance data can be collected continuously before it must either be deleted or transmitted.	189
4.13	The calibration circuit used for the laboratory validation of the wearable device, reminiscent of those shown in Figure 3.5. Two potentiometers, R_1 and R_2 , and a capacitor, C , are meant to act as the extracellular and intracellular components.	193

4.14	Measurement results of resistance and reactance from the wearable device configured for both bipolar (dashed line) and tetrapolar (marked with circles). Low, medium, and high impedance values are shown at the top, middle, and bottom, respectively.	195
4.15	A typical resistance and reactance plot for the equivalent circuit measurements before (dashed lines) and after (solid lines) calibration for the bipolar (blue) and tetrapolar (red) configurations. The predicted response of the circuit (the one shown in the middle of Figure 4.14) is shown in black. In some areas it may be hard to see either the bipolar or tetrapolar response because of how well they align with predictions.	197
4.16	Results of the wearable device on three human subjects. At the 10 second mark of each graph the reconfigurable front end switched from the bipolar to the tetrapolar arrangement. (Moderate switch bouncing was observed and removed from this graph for the sake of clarity.) The top subject shows a response dominated by the respiratory signal with little evidence of a cardiac component. The bottom subject shows an opposite response, one where the cardiac signal dominates over respiration. In the middle of the two, both figuratively and literally, is a subject with both a cardiac and respiratory response. In all cases, the wearable device's results agreed well with those found via the EBI100C.	198

LIST OF TABLES

1.1	Diameter of blood vessels, the velocity of blood flowing within them, and their corresponding Reynolds and Womersley numbers. From this the wide range of diameters, velocities, Reynolds and Womersley numbers can be appreciated. All values were calculated assuming a heart beat of 75 beats per minute, a blood density of 1050 kg/m^3 (at body temperature), and a blood viscosity of 4 centipoise. Within the body, each of these parameters exist within a range, but for an order of magnitude assessment of the situation, this table should suffice.	22
2.1	Acoustic properties of various biological materials. Listed in this table are the density, bulk modulus, speed of sound, attenuation coefficient, and characteristic acoustic impedance for many materials that are relevant to this topic. Where applicable the mean value of a parameter for a material is accompanied by a standard deviation in parenthesis. Values are taken from Hasgall et al.	62
2.2	The averages, standard deviations, ranges, and intraclass correlation coefficients for the maximum IVC diameter, the minimum IVC diameter, and the caval index as measured by physicians and the proposed algorithm.	116
3.1	Dielectric properties of a range of biological tissues. The parameters fit to the parameterized Cole-Cole summation as seen in equation 3.23. Here, k and M represent multiplying the value by 10^3 and 10^6 , respectively. The materials displayed here are the same as those reported in 2.1.	131
3.2	A summary of relevant aspects of current volume status assessment techniques. For a fuller examination of these approaches, see the section titled <i>Monitoring</i> in Chapter 1. (CVP - central venous pressure, PAOP - pulmonary artery occlusion pressure, RVEDV - right ventricle end-diastolic volume, LVEDA - left ventricle end-diastolic area, IVC - inferior vena cava, BIVA - bioimpedance vector analysis, SVV - stroke volume variation, PPV - pulse pressure variation, dIVC - change in inferior vena cava diameter, NICOM - non-invasive cardiac output monitoring, typically referring to bioreactance here, I2VE - impedimetric intravascular volume evaluation.)	145

4.1	A list of a few biomedical signals currently able to be measured through wearable electronics. Listed are the signal, a few sensors types currently used to measure them, and the signal source. Dividing the five vitals at the top of the list and several other practical signals is my own I2VE, whose status I believe ought to be treated as a measure of a vital sign (volume status) whose time shall come.	172
4.2	List of components used for the wearable electronic device. (Unless otherwise stated all passive surface-mounted devices are 0603 package sizes.)	190
4.3	Measurement errors in resistance and reactance for each front end configuration without calibration, with a first order calibration, and with a second order calibration. Results are presented as an average with a 95% interval in brackets.	196

ABSTRACT

Many medical conditions require health care providers to make rapid and accurate assessments of a patient's hemodynamics to diagnose and guide treatment. Of the parameters that encompass hemodynamic evaluation (such as heart rate and blood pressure), one remains woefully measured: intravascular volume status. A measure of how much blood a subject has with which to perfuse oxygen, deliver nutrients, and remove wastes, volume status lays the foundation of many types of fluid based treatments for the acutely and critically ill in emergency rooms and intensive care units to the chronically ill of cardiovascular and dialysis centers. Yet for all its physiological importance, many current assessments of volume status and responsiveness have a predictive power around the level of a coin flip. This work seeks to remedy this dire situation by improving a technique already utilized at the bedside (ultrasound measurement of inferior vena cava collapsibility) and introducing another that could be used through all echelons of care (impedimetric intravascular volume evaluation). Specifically, three major accomplishments were achieved: (1) an open-source software toolbox was developed to help clinicians and researchers rigorously evaluate vessel strain, stress, and wall shear; (2) a novel bioimpedance technique was demonstrated in several patient populations; and (3) a wearable electronic device was created to continuously measure a subject's volume status. Such a multimodal approach to hemodynamic monitoring serves the needs of both physicians and their patients across many clinical environments and ought to improve outcomes.

CHAPTER 1

An introduction to hemodynamics

For the heart, life is easy. It beats for as long as it can, then it stops.

Karl Ove Knausgaard, *My Struggle*

At its most basic level the cardiovascular system is a collection of tubes filled with fluids and a pump to move them. The whole point of this setup is to deliver volumes of blood to regions of the body primarily for the purposes of transporting dissolved gases, nutrients, hormones, and metabolic wastes to places they are needed from places they are not, and regulating the pH balance and ion concentrations of interstitial fluids to ensure a consistent biochemistry throughout the body. Blood coursing through the body performs many other functions including stabilizing body temperatures, defending against the introduction of toxins and the invasion of pathogens, and restricting fluid losses at sites of injury, but for the investigation reported herein, our attention is best focused at on the transportation and the regulation of fluids. The pump and the tubes shall concern us here.

Even for such a simple system, many problems can arise. Problems can arise with the heart (the pump) including coronary artery disease, arrhythmia, heart valve disease, congenital heart disease, cardiomyopathies, pericarditis, aorta disease, Marfan syndrome, and myocardial infarction. Problems can arise with the vasculature (the tubes) including peripheral vascular disease (one of the fast growing causes of mortality world wide [1], aneurysm, renal artery disease, Raynauds syndrome, Buegers disease, blood clotting disorders, and lymphedema. Problems can arise with the fluids themselves in the form of anemia, lymphoma, and hemophilia, how the fluids are regulated as is the case in chronic kidney disease, diabetes, and end stage renal disease, and how the fluids are distributed from hypertension to edema. These are but a small fraction of all possible problems, but together these problems represent tens of millions of lives lost and hundreds of billions of dollars spent annually [1]. Such costs — in money, in life, in suffering — are high, and growing higher.

Treatments for these and many other conditions are predicated on clinicians abilities to diagnose the presence and extent of the conditions themselves. Without a semblance of what is wrong, one would be at great pains to know how to make it right. Thus stands the entire edifice of clinical decision-making on the foundation of tools to measure the regularity and irregularity of patient states. Disorders, diseases, and pathologies manifest themselves in many and multifaceted ways, affecting the bodys anatomy, physiology, chemistry, and functions. The detection of such anomalous behavior lies at the heart of clinical monitoring.

Herein I present one solution within a subset of clinical monitoring. It is a mixed modal approach linking the physiology detecting methods of bioimpedance with the anatomy imaging methods of ultrasound. In their combination is a hybrid method of evaluating parts of the human body in a way heretofore unavailable to clinicians and researchers. This mixed modal technique allowed for the measurement of long-term changes to anatomy (as seen in a clinical study presented in Chapter 2), rapid, short-term dynamic changes in anatomy and physiology (as seen in the clinical trials presented in Chapters 2 and 3), and finally the development of a wearable clinical monitor utilizing clinical results to continuously monitor these structural-functional changes (as discussed in Chapter 4). The whole of this work is intended to demonstrate the relatedness of the signals derivable from electrical impedance (bioimpedance) and acoustical impedance (ultrasonography) based techniques. Though the signals correspond in many ways, such as the results determined from the ultrasound measurement of the inferior vena cava and the measure of electrical impedance in the limb, they differ greatly in the ways they are collected, examined, and interpreted. The goal of this work is to demonstrate a pathway towards the ultimate hybridization of these techniques.

The path taken for this demonstration began with the development and implementation of separate signal and image processing techniques for both bioimpedance and ultrasound. Algorithms and software were created in parallel for each modality to independently verify their accuracy before then synthesizing their results.

I first designed an ultrasound imaging toolbox to work with and analyze post-clinical images in the standard Digital Imaging and Communications in Medicine (DICOM) format. I chose this approach given the need to test this work in many different clinical environments where few researchers and clinicians have access to the raw underlying radiofrequency data of their ultrasound machines. The other tangible benefit to such an approach is in the creation of a software toolbox that can be used by researchers around the world to identify and measure parameters of interest freely and without the need to modify their existing imaging capabilities. The specifications and validation of this toolbox will be dis-

cussed at length in Chapter 2 and it will suffice to state here that the software was created to make possible the measurement of wall distensibility, local and global tissue strain, and the shear rate of blood across vessel walls three of the most prominent and significant dynamically varying aspects of the vasculature measurable by ultrasound.

Once mature, the ultrasound software was used to compare a novel bioimpedance algorithm to assess dynamic changes in intravascular volume against a standard ultrasound procedure. These techniques were tested in three separate patient populations: normal volunteers and not severely ill or injured emergency department patients; patients undergoing hemodialysis; and mechanically ventilated patients in an intensive care unit setting. The design, justification, and results of these experiments are presented in Chapter 3.

Once validated, a small, wearable version of the bioimpedance monitoring system was developed. An electronic circuit was designed, developed, constructed, and tested both in simulated environments and on healthy volunteers. This work represents the majority of Chapter 4. Future work on additional patient populations is proposed for further validation and prognostic use. By this way, I hope to start down the path of measuring the physiology of subjects in any and all environments.

To understand any measurement technique, one must first understand the system being measured. Only in such a manner can the signals derived from the measurement be understood. That is to say, without context a signal is without meaning. For us to impart meaning on the work presented here, let us first understand the environment in which it is collected. Broadly speaking this work was conducted in two general environments: the pre-clinical and the clinical. The preclinical work consisted of computation fluid dynamic simulation of blood flowing through peripheral vasculature, particle image velocimetric validation of such simulations as a standard for our developed software algorithms, and in the design and calibration of a wearable monitoring device. The software and hardware components developed were then tested in one of four clinics at the University of Michigan depending on the study: the transplant unit (for follow up ultrasound measurements of fistula), the emergency medicine department (for the initial ultrasound correlation with my proposed bioimpedance technique), the in-hospital dialysis unit (to track large volume shifts via ultrasound and impedance), and the intensive care unit (to validate converse pressure relationships). In so doing, I have endeavored to move these techniques and this technology from the bench to the bedside as quickly as possible, validating the two modes in each environment. For all of this, ultimately generating each and every signal examined here from the strain and wall shear measurements in the fistula maturation study (Chapter 2) to the volume status monitoring in the dialysis clinic (Chapter 3) were the hearts of each and

every patient.¹ Let us begin then at the heart of the matter.

1.1 The heart

As Knausgaard reminds us in the introductory quote for this chapter, life is easy for the heart: It beats for as long as it can, then it stops. In fact, a person may expect to have their heart beat more than two billion times over the course of an average life span. But within each beat is a complex cycle of contraction and relaxation, known as systole and diastole, respectively, allowing the heart to push blood from its four chambers and to receive blood in preparation for the next cycle. The heart's sole obligation to us and to our body is to pump blood along two vascular pathways. The first of these pathways is organized into the pulmonary circuit, a network of vessels that carry blood to and from the gas exchanging surfaces of the lungs. The second pathway is called the systemic circuit and it is along this path that blood is transported to the rest of the body. Both of these circuits begin and end at the heart as blood courses first from one circuit then to the other then back again in sequence. A schematic representation of these pathways and their relationship to the heart can be seen in Figure 1.1.

Of the four chambers of the heart, the two on top, called atria, can be thought of as blood receivers, and the two on bottom, called ventricles, can be thought of as the blood pushers. Starting from a person's right atrium, blood is received by the heart from the systemic circuit, then pushed to the right ventricle. From the right ventricle blood is pushed into the pulmonary circuit, the blood becomes oxygenated, and returns to the heart via the left atrium. The left atrium collects the blood and empties it into the left ventricle, which then pumps the blood into the systemic circuit, and the whole process begins anew.² (For graphical representation of this flow pattern, please review Figure 1.2) As the heart beats, first the atria contract, emptying their contents into their respective ventricles, then the ventricles contract. Normally this ventricular contraction occurs in both ventricles simultaneously and with equal volumes of blood being ejected into the pulmonary and systemic circuits. This, however, is not always the case.

Blood leaving the left ventricle passes into the ascending aorta (the largest vessel in the human body) through a valve, a thick membrane designed to open and close with respect to pressure, in this case, the aortic valve, also known as the aortic semilunar valve. Similar valves exist between the atria and the ventricles in the heart. On the right side of the heart

¹The kidneys and the lungs also played a significant role in the two studies named. The author hopes to be forgiven for a poetic flourish by correcting the record here.

²A more detailed explanation of the blood flow and pressures occurring during this process are discussed in the next section




Figure 1.1: The circuits of the circulatory system. In human beings the pulmonary circuit begins at the right ventricle where blood is ejected to the lungs and ends at the left atrium where oxygenated blood enters the heart. The system circuits begins with blood being ejected from the left ventricle and with it returning through the right atrium. Image is from [http://msjensen.cehd.umn.edu/ Webanatomy/imagetatabase/Cardiovascular/circulation.tif](http://msjensen.cehd.umn.edu/Webanatomy/imagetatabase/Cardiovascular/circulation.tif) and is used for educational purposes.




Figure 1.2: Basic anatomy of the heart demonstrating the relationships between the atria, the ventricles, and the vessels what feed into them. The flow of blood within the heart is indicated with arrows. Modified from "Diagram of the human heart (cropped)" licensed under CC-BY-SA-3.0 via Wikimedia Commons.

lies a set of three fibrous flaps that constitute the right atrioventricular valve, also known as the tricuspid valve. On the left side of the heart is a pair of fibrous flaps referred to variously as the left arterioventricular valve, the bicuspid valve (as distinct from the right tricuspid valve), and the mitral valve. Each of these valves is meant to serve the same purpose: to prevent blood from flowing in an undesired direction. Said conversely, it is meant to keep blood flowing in the correct direction through the heart and into the vascular circuits. To understand why we must first consider a basic aspect of fluids and their movements: fluids, be they water through a pipe or blood through the cardiovascular system, move down pressure gradients, moving from areas of higher pressure to ones of lower pressure. This basic principle allows blood to flow from the high-pressured aorta down to the arteries, into the arterial and ventricles, up the veins, back into the heart. (For a basic representation of the pressures throughout the body see 1.6) In the heart itself, the pressure in each chamber rises during systole, falls in diastole, forcing blood to move around and through and into the heart. The valves of the heart, between the chambers and at the aorta, ensure blood flows in the correct direction, by opening to positive pressure gradients in one direction and closing to positive pressure gradients in the opposite direction. In this way, the valves attempt to establish a consistent unidirectional pressure gradient throughout the entire circulatory system.

To maintain correct pressure relationships throughout the body, the cardiovascular reacts quickly to changes in volume and velocity via mechanotransduction occurring along and within vessel walls. The specifics of this maintenance will be elaborated in the coming sections on the vasculature itself.

For now, let us begin by imagining the start of an average idealized cardiac cycle with a resting heart rate of about 75 beats per minute, a systolic blood pressure of 120 mmHg, and a diastolic blood pressure of 80 mmHg. All four chambers are relaxed and the ventricles are partially filled with blood. First, the atria contract, increasing atrial pressure, causing the ventricles to fill completely with blood. Because atrial pressure exceeds venous pressure blood does not flow into the atria during this time. At the end of atrial systole, each ventricle contains the maximum amount of blood that it will hold during a cycle, a quantity known as end diastolic volume (EDV). Next the atria relax, entering a period of atrial diastole that lasts until the next cardiac cycle. Starting at the same time as this atrial diastole is ventricular systole, the period of ventricle contraction. During this stage ventricles begin by contracting isometrically, inducing a pressure rise ultimately reaching a point of isovolumetric contraction where all heart valves are closed, the volumes of the ventricles remain constant, and ventricular pressures rise above the atrial. Once pressure in the ventricles exceeds that in the arterial trunks, the semilunar valves of the heart open and blood flows

into the pulmonary and aortic trunks, marking the beginning of ventricular ejection. During ventricular ejection, each ventricle will eject some volume of blood (typically 70-80 mL, but can range anywhere from about 60-100 mL), known as the stroke volume of the heart. The ratio of this stroke volume to the end-diastolic volume is known as the ejection fraction and varies in response to the changing demands on the heart (to be discussed at a later point). At this point the ventricles contract isotonicly. Having reached a peak, ventricular pressures decline rapidly and blood in the aorta and pulmonary trunks starts to flow back toward the ventricles, causing the semilunar valves to close. As these valves close, there is a small temporary rise in pressure as the elastic arterial walls recoil, producing a distinct valley pressure tracing known as the aortic notch. The amount of blood that remains in the ventricle when the valve closes is known as the end-systolic volume (ESV). The heart then enters a period of ventricular diastole for the final stage of the cardiac cycle. The ventricles enter a period of isovolumetric relaxation until ventricular pressures fall below those of their atrial counterparts and the AV valves are forced open. The pressures in the ventricles are far below that of major veins, allowing blood to pour into the relaxed atria and ventricles, passively filling the chambers of the heart before the whole process is repeated.

The consistent cycling of the heart produces many possible signals to measure. From the pressures within the atria, ventricles, and the vessel feeding them and the volume of blood present in the heart at any given time to the electrical potentials caused by the heart's shifting dipole and even the sounds it makes, this process generates each of the signals that are of interest to this work. It is therefore instructive to note here the relationship between these parameters as a function of the cardiac cycle. One method of doing so was developed by Carl J. Wiggers, a cardiovascular physiologist working at (Case) Western Reserve University in the early half of the twentieth century, where each of these signals is plotted graphically in parallel across time. A representative example of this Wiggers diagram can be seen in Figure 1.3, from which one can appreciate the synchronicity of the electrical, volumetric, and pressure changes constantly occurring across the heart. Special attention should be paid to the pressure and volume shifts as it is from these that much of this work builds from.

Ultimately of greatest concern in the pumping of the heart is how much blood is expelled from the ventricles for each beat. This quantity, known as the stroke volume, is equivalent to the end-systolic volume (the amount of blood remaining in each ventricle at the end of systole) subtracted from the end-diastolic volume (the volume of blood in each ventricle at the end of diastole). For a single cardiac cycle, stroke volume may be the single most important factor to examine because this is fundamentally a quantification of the




Figure 1.3: Pressure and volume relationships during systole and diastole. Modified from "Wiggers Diagram" by DanielChangMD revised original work of DestinyQx; Redrawn as SVG by xavax - Wikimedia Commons, File:Wiggers Diagram.png.. Licensed under CC BY-SA 2.5 via Wikimedia Commons

hearts single beat efficiency.

To gauge a hearts efficiency over a longer period of time than a single cardiac cycle, physicians will often refer to a quantity known as the cardiac output, the amount of blood pumped by the left ventricle per some unit time. This value, a measure of ventricular efficiency over time, is often calculated by multiplying the heart rate (in beats per minute) by the stroke volume (in mL per beat) and generally falls within the range of 4-8 L/min [2]. Cardiac output is essentially the volumetric flow rate of blood through peripheral tissues, indicative of what the heart and body perceive to be necessary to maintain homeostasis. To keep the body stable throughout many different activity levels (from resting to running) and health states (from picture perfect to pathological), the heart adjusts its cardiac output to balance the needs of the systems. A schematic representation of the factors affecting cardiac output are shown in Figure 1.4.

Broadly speaking there are two ways the heart is able to adjust its output, either through

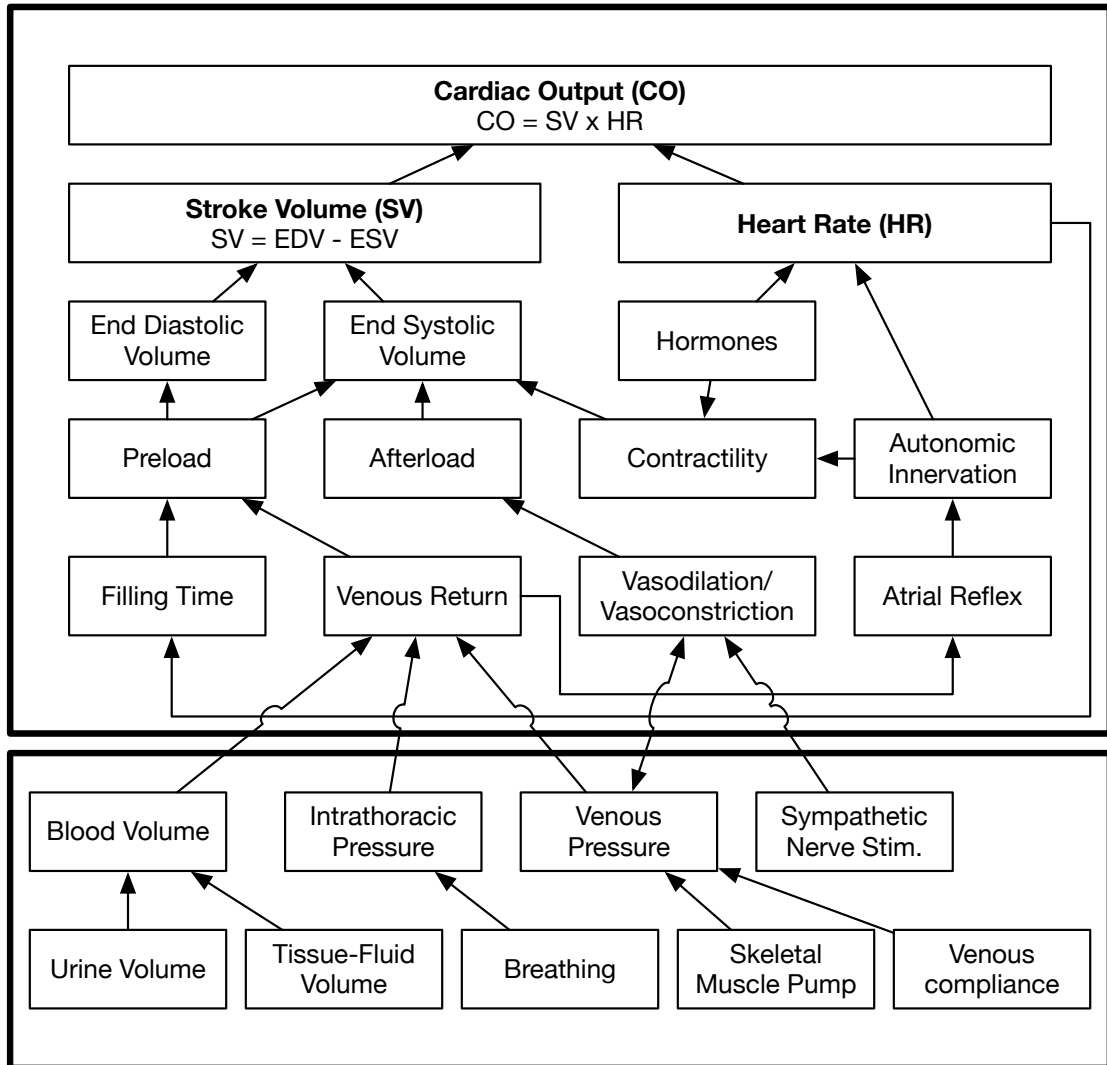


Figure 1.4: Schematic representation of the main factors affecting cardiac output. Cardiac output, typically represented as the product of stroke volume and heart rate, is a measure of the volumetric flow rate of blood pumped out of the ventricles to the pulmonary and systemic circuits. It is affected and affects many hemodynamic parameters and its relationship to the larger fluid volume management capabilities of the body must be understood in context.

the heart rate or the stroke volume. Though an area of intense research and clinical relevance, the electrochemical generation of heart contractions and therefore one of the principle components of the heart rate will not be discussed here at any great length. Going still further, in our haste we will also forgo discussions of cardiac reflexes and autonomic tone on the heart rate. An exception will exist in the case of atrial reflex, which is the hearts response to the walls of the right atrium stretching in response to increased venous return. This reflex, also known as the Bainbridge reflex, is one case where volume is mechanically transduced by the system to cause a rate change, rather than another volume change. That both types of change, be they rate or volume, indirectly affect pressure gradients and flows should not go unnoticed. These factors are intrinsically bound in the bodys feedback-control system of regulation and more often than not, all else being equal, one will affect the other in some way. It might be intuitive for some readers to consider these two parameters, rate and change, as equivalent to a frequency and amplitude, respectively, of some ideal signal the heart is attempting to create. However, our concerns for this work will lie primarily with the effects of volumes on hemodynamic parameters. Therefore, we will concern ourselves primarily with those factors affecting stroke volume.

Stroke volume, as previously mentioned, is the difference between the end-diastolic volume and the end-systolic volume. How these volumes change in response to a myriad of conditions has been the subject of decades of intense medical research. For now a brief overview of the factors affecting each should suffice.

End-diastolic volume is affected by two main factors: the filling time and venous return. The filling time is simply the time it takes to complete ventricular diastole, itself a function of heart rate. As the heart rate increases, filling time decreases. To go back to my previously hinted at metaphor, this would be equivalent to adjusting the frequency of the ideal cardiac signal. To adjust the amplitude of this signal, one would turn to venous return. Strictly speaking venous return is merely the rate of blood flow returning to the heart during this diastolic period and it changes in response to cardiac output, blood volume, patterns of peripheral circulation, and intrathoracic pressures [3]. By increasing the volume of blood returning to the heart, and holding all other parameters constant, one could increase the volume of blood exiting the heart. Increasing venous return tends to increase end-diastolic volume, which can lead to an increase in stroke volume.

Augmenting filling time and venous return changes a parameter of the heart itself known as the cardiac preload. Depending on the parlance, this preload refers to either the end diastolic pressure responsible for or the degree to which the ventricles of the heart stretch (alter their greatest geometric dimensions) under variable physiologic demands [4, 5]. Put simply, preload is the measure of myocardial stretching in the ventricles during

diastole. This amount varies with the demands of the heart and is directly proportional to the end-diastolic volume itself. Increasing preload increases the end-diastolic volume. All else being equal, stroke volume should increase with preload. And within normal physiological limits this is generally true. First proposed by Ernest Starling and Otto Frank, this relationship of more in equals more out is essence of the Frank-Starling law of the heart [5]. How end-diastolic volume and right atrial pressure affect stroke volume and cardiac output (and therefore venous return) is shown in Figure 1.5. There are of course instances where the Frank-Starling law does not hold [6, 7], but for the majority of cases, it is deemed reliable enough for clinical decision making.

As popularized by Guyton, the operation of the heart at any given time is determined by the intersection of the Frank-Starling cardiac output curve with the blood volume return curve of the patient as seen in Figure 1.5. There are many factors affecting the shapes, slopes, and intercepts of these two curves, but generally speaking contractility and blood volume play a significant role in their behavior. As the volume of blood is increased, we can see that preload and output increase in turn and is the foundation for fluid replacement therapy. The converse is true and demonstrates the dangerous consequences of hemorrhage and massive blood loss. A similar direct relationship can be seen with increased contractility increasing the operating point of the heart and decreased contractility decreasing the operating point. Add the myriad other factors that can alter these curves' behaviors in conjunction with the fact that similar values of preload can give wildly varying values for cardiac output and some the problems of hemodynamic monitoring become apparent.

End-systolic volume, the amount of blood remaining in the ventricle at the end of ventricular systole, is subjected to three main factors. The first is preload which we have discussed previously. How this affects the ending volume is obvious: the amount that is left over is necessarily a function of the amount initially available. The second major factor is the contractility of the ventricle. Contractility, represented by the force required for contraction at a given preload, represents the intrinsic ability of the myocardium to contract. Under normal circumstances, the ability to produce force during contractions is altered by autonomic innervation and many circulating hormones, each acting on the degree to which myosin and actin in these muscles are bound, a relationship itself determined by the calcium ion concentrations in the cardiac cells. Both sympathetic and parasympathetic stimulation affect the contractility of the heart and can be thought of as either a positive inotropic action, if it increases contractility, or a negative inotropic action, if it decreases contractility. As contractility increases, end-systolic volume decreases and vice versa. The final factor affecting end-systolic volume is the afterload, the amount of tension the contracting ventricles must produce to open the semilunar valve and eject blood. Afterload relates directly to

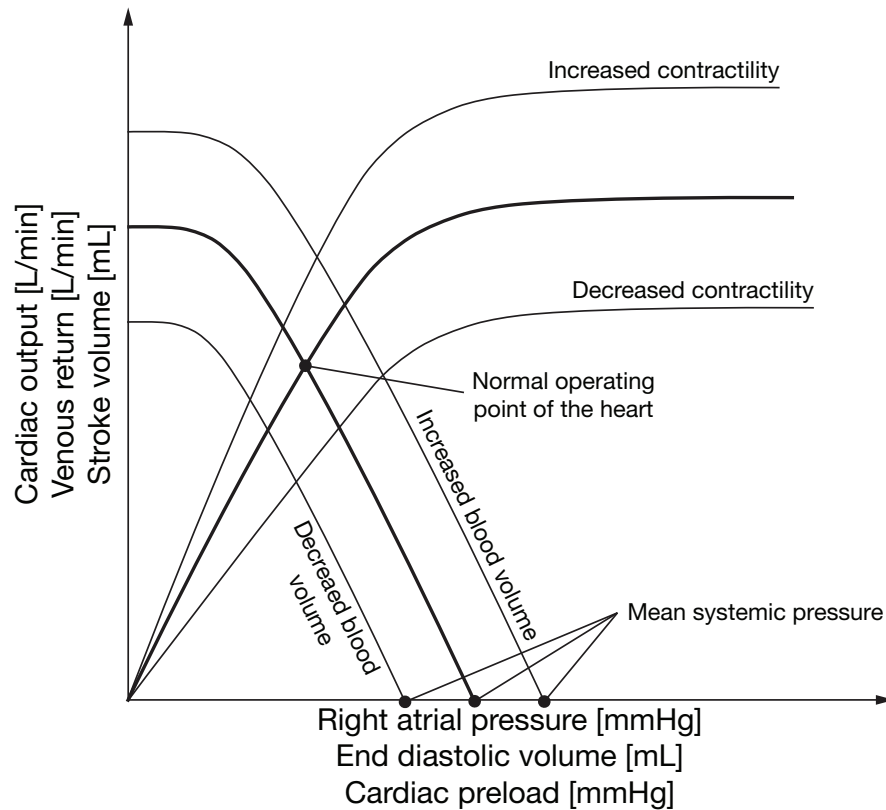


Figure 1.5: The cardiac function curve as predicted by the Frank-Starling law and its coupling with venous return. The intersection of the two curves represents the operating point of the heart. Changes to either the cardiac output or blood volume curve affect how much blood can be circulating and in what way the heart circulates that blood.

the period of isovolumetric contraction (the greater the afterload, the longer the period of contraction) and somewhat inversely to the duration of ventricular ejection (the greater the afterload, the shorter the period of ejection). As such, increased afterload tends to increase end-systolic volume and thereby decrease stroke volume.

Multiplying the volume ejected by each pump of the heart by the rate at which each pump occurs, yields the aforementioned cardiac output in terms of a volumetric flow rate. It is this volumetric flow that will be pushed to every portion of the body, providing nutrients, giving oxygen, and removing waste from every single living cell of our bodies. How this flow is created by the heart is crucial to understanding how the vasculature operates: flow from the heart is the primary input to the blood vessels system of transportation. With this basic understanding of how the heart operates, we can now proceed through the pump into the pipes, where blood volumes flow, perfuse, and are returned to whence they came. We can now begin to understand the role of the vasculature.

1.2 The vasculature

Using the heart again as our starting point, we can classify vessels by whether they carry blood away from the heart, the arterial side, or return blood back to the heart, the venous side. Specifically, blood is pumped from the heart, into the arteries (also known as efferent vessels), transported to peripheral tissues via smaller arteries known as arterioles, and through the capillaries where diffusion occurs between blood and interstitial fluids. After the capillaries, blood flows into the venules and finally the veins (also known as afferent vessels) on its way back to the heart. This relationship is true whether the blood is being carried through the pulmonary circuit on its way to and from the gas exchanging surfaces of the lungs or its being carried along the systemic circuit throughout the rest of the body.

As blood courses through the vasculature, the vessel diameters tend to go from large near the heart (the aorta, the largest artery, is typically around 2.5 cm, while the vena cava, the largest vein, can be up to 3 cm) to their smallest point within the capillaries (this relationship can be observed in Figure 1.6(a)). This decrease in individual vessel size is accompanied by an increased number of vessels branching further along the pulmonary and systemic circuits. The effect of this vessel branching causes the total cross-sectional area of the vessels to increase as blood is pumped to and from the capillary beds (Figure 1.6(b)). As previously mentioned, blood will only flow down a pressure gradient and this true throughout the vasculature where the average blood pressure continues to decrease from its peak, just exiting the aorta, to its lowest point, returning through the vena cava (as seen in Figure 1.6(c)). Blood flow, being proportional to the ratio of pressure to cross-sectional area, decreases through the arterial system, reaches a minimum necessary for sufficient exchange and perfusion to occur throughout the capillary beds, then begins to rise throughout the venous system as blood returns to the heart (see Figure 1.6(d)).

Anatomically, arteries differ from veins in a number of significant ways. Arteries typically have thicker, more muscular walls than veins enabling arteries to resist the large pressure change generated by the heart as it forces blood through the circulatory system by passively responding to changes in pressure through elastic strain. Given their contractility, arteries may also actively change their diameter (via the sympathetic autonomic nervous system) by either decreasing their size, a process known as vasoconstriction, or by increasing it, through vasodilation. This active control affects a number of different parameters including peripheral blood pressure, capillary blood flow, and afterload of the heart. Veins, in contrast, have thinner walls, are more elastic, often tend to have larger diameters than their arterial counterparts, and tend not to be able to tolerate as much force. The venous system also contains valves to prevent blood from flowing backward.

Between the two systems lie the capillaries. These very small, very delicate vessels are woven through us, feeding our muscles, surrounding our connective tissues, allowing our skin to breathe. Bundled in networks called capillary beds, it is only through capillaries that blood can exchange what it needs (water, solutes, gases, etc.) with interstitial tissues and fluids. To make this possible, the vast majority of the walls of these vessels are so thin (5 m) and so small (8 m, around the size of a single red blood cell) that diffusion can occur across them. Capillaries of this sort that allow for continuous passive diffusion are known as continuous capillaries and they are located in all tissues except for epithelia and cartilage. In contrast, fenestrated capillaries contain small pores in their endothelial lining that allow for rapid exchange of fluids and small peptides. Fenestrated capillaries are located in a number of areas, including hormone secreting regions (the hypothalamus, the pituitary gland, and the thyroid), absorptive areas (the intestines), and in areas of filtration, most notably the kidneys. Once the blood in these beds has done its work it must come back to the heart. From all of the branches the vessels have broken into, they must recombine. This joining of two (or more) vessels is known as an anastomosis. Anastomoses serve many functions throughout the body, acting as backup routes for blood flow (as is the case in arterial anastomosis) and as connections from the arterioles to venules (as is the case in arteriovenous anastomoses). The role of artificial anastomoses in clinical practice will be discussed at length later in and will lie at the foundation of the work conducted in Chapter 2.

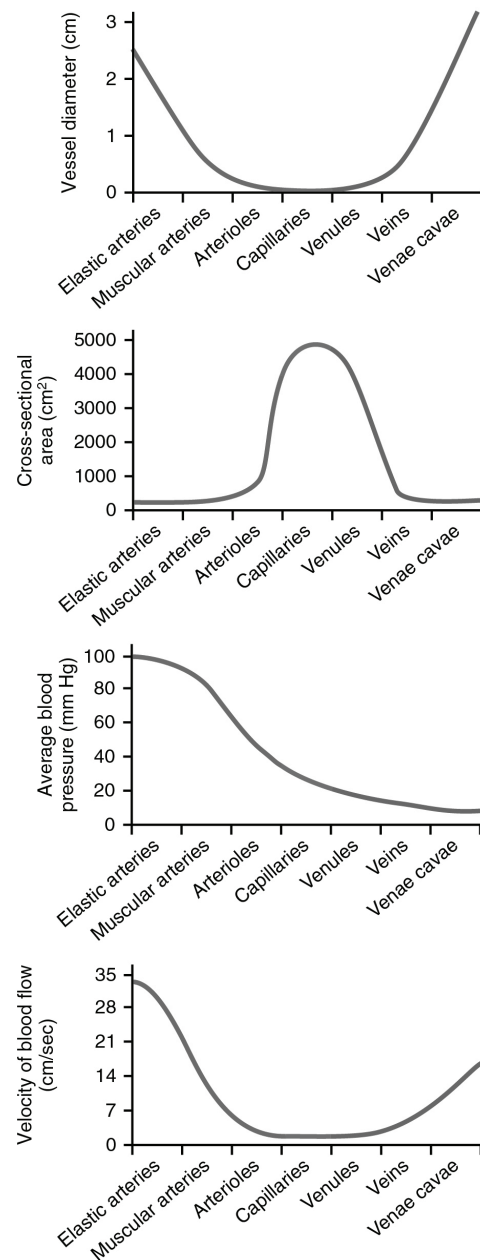


Figure 1.6: The relationship between the (a) diameter, (b) cross-sectional area, (c) average blood pressure, and (d) velocity of blood with respect to their vessel type in the cardiovascular system.

Blood does not reside equally across the body. Typically, about one third of total blood volume (1.5 L) is contained in the heart, arteries, and capillaries of the pulmonary and systemic circuits. The other two thirds (3.5 L) is retained in the venous system with a large portion of that amount (1 L) circulating within the extensive venous networks of the liver, bone marrow, and the skin. Blood is retained in this low-pressure to serve as a reservoir for many compensatory mechanisms necessary for homeostasis. Because of their generally thinner walls and lower proportion of smooth muscle, veins are able to distend (expand) far more than arteries, accommodating large changes in blood volume. In this sense, veins have a greater capacitance — the relationship between a vessels volume and the pressure at which it can hold that volume — than arteries. That is, because it takes less pressure to expand veins than it does arteries, they have a great capacitance. Herein, to avoid confusion with the electrical term, I will refer to this quantity—represented by the ratio of the change in volume to the change in pressure $\Delta V/\Delta P$ —as the compliance of the vessel. The compliance of arteries and veins will play a key role in each of the clinical investigations.

The venous system is also able to contract to compensatory for changes throughout the cardiovascular system. This contraction, venoconstriction, reduces the amount of blood in the venous system to attempt to increase the volume within the arterial system and the capillaries. With such a large reservoir (and such a fast ability to respond), the venous system can help maintain arterial blood volume at near-normal levels despite potentially large changes in volume blood volume. In cases of large volume loss, such as in hemorrhage and overly aggressive hemodialysis treatment, the extreme venous response can have the adverse consequence of increasing vascular resistance, making it difficult for blood to return to the heart and, aggravating already delicate situations.

It would be fruitful now to remember that pressures and volumes throughout the cardiovascular system are by their nature dynamic. The nature of the cardiac cycle itself ensures that there are peak pressures, shifting volumes, and constantly changing flows. From a control theory perspective, the body constantly undergoes excitation from these varying parameters and attempts to reconcile their effects into a cohesive homeostasis. How the bodys systems maintain or fail to maintain the feedback loops necessary to achieve these ends is a crucial component of diagnosing pathological states and, in the end, for treating or improving them. The work outlined here is meant to be a tool to help diagnose such states, by estimating volume status, a key parameter of hemodynamic stability. But before we can understand any sort of signal from any sort of system, we must first understand its context, its environment, and the way in which it is collected. Therefore, it will be necessary to first understand the manner in which blood flows through the vessels, current means of mea-

asuring certain hemodynamic principles, and finally how each of the signals are modified in normal and abnormal states.

1.3 Blood flow

Simplifying the situation a bit, the flow of blood (and of any fluid really) can be considered a function of the ratio of the pressure gradient to the systems resistance to that flow. This can be written quasi-mathematically by saying F is proportional to $\Delta P/R$, where F is the flow, ΔP is the pressure gradient, and R is the resistance. Defining it in such a way allows us to grasp the fundamental basis of bulk flow: there are forces attempting to drive a fluid in one direction and there are forces attempting to prevent the fluid from flowing. Even this basic conceptual framework is enough to understand many significant portions of the circulatory system.

The entire cardiovascular pressure gradient starts and ends at the heart. Here pressures reach their peak and their valley (averaged over a full cardiac cycle). Just exiting the heart, the aorta and the elastic arteries experience very large pressures and very large pressure changes (see the difference in systolic and diastolic curves in Figure 1.3). Pressure in the arteries, often referred to simply as blood pressure, varies tremendously (see Figure 1.6), and as such must be taken from a consistent location to make sense of any time related trends. Passing through the arterioles into the capillaries, the pressure known as the capillary hydrostatic pressure continues to drop in magnitude on its way to the venous system where the cardiovascular reaches its minimum values venous pressure and, ultimately, right atrial pressure. Though in the course of normal cardiovascular function pressures will also decrease to push blood ever onward from (and back to) the heart, we must remember that it is the pressure gradients that affect flows, and not pressures themselves. (A brief examination of these relationships for real physiological flows can be seen in Section 2.4.1.6.)

A pressure gradient is merely the potential for fluid to flow, the resistance to that potential will determine in what manner and to what degree that fluid flows. In the human body, many factors affect the resistance to flow. Chief among these factors are those affecting vessel geometry (length, diameter, curvature, tortuosity, etc.) and blood viscosity (itself affected by temperature, hematocrit level, etc.). To see in what ways each of these parameters affects fluid flows, let us turn our attention to basic fluid modeling.

To model a fluid fully one may begin with the Navier-Stokes equations. Their derivation and proof will not be considered here and their role here is merely as a tool to explain the motion of blood in a few very specific contexts. Given this environment, we will begin by

expressing the equations in cylindrical coordinates in full.

$$\begin{aligned} \rho \left(\frac{\partial v_r}{\partial t} + v_r \frac{\partial v_r}{\partial r} + \frac{v_\theta}{r} \frac{\partial v_r}{\partial \theta} - \frac{v_\theta^2}{r} + v_z \frac{\partial v_r}{\partial z} \right) \\ = \rho g_r - \frac{\partial p}{\partial r} + \mu \left[\frac{\partial}{\partial r} \left(\frac{1}{r} \frac{\partial}{\partial r} (r v_r) \right) + \frac{1}{r^2} \frac{\partial^2 v_r}{\partial \theta^2} - \frac{2}{r^2} \frac{\partial v_\theta}{\partial \theta} + \frac{\partial^2 v_r}{\partial z^2} \right] \end{aligned} \quad (1.1)$$

$$\begin{aligned} \rho \left(\frac{\partial v_z}{\partial t} + v_r \frac{\partial v_z}{\partial r} + \frac{v_\theta}{r} \frac{\partial v_z}{\partial \theta} + v_z \frac{\partial v_z}{\partial z} \right) \\ = \rho g_z - \frac{\partial p}{\partial z} + \mu \left[\frac{1}{r} \frac{\partial}{\partial r} \left(r \frac{\partial v_z}{\partial r} \right) + \frac{1}{r^2} \frac{\partial^2 v_z}{\partial \theta^2} + \frac{\partial^2 v_z}{\partial z^2} \right] \end{aligned} \quad (1.2)$$

$$\begin{aligned} \rho \left(\frac{\partial v_\theta}{\partial t} + v_r \frac{\partial v_\theta}{\partial r} + \frac{v_\theta}{r} \frac{\partial v_\theta}{\partial \theta} + \frac{v_r v_\theta}{r} + v_z \frac{\partial v_\theta}{\partial z} \right) \\ = \rho g_\theta - \frac{\partial p}{\partial \theta} + \mu \left[\frac{\partial}{\partial r} \left(\frac{1}{r} \frac{\partial}{\partial r} (r v_\theta) \right) + \frac{1}{r^2} \frac{\partial^2}{\partial \theta^2} - \frac{2}{r^2} \frac{\partial v_r}{\partial \theta} + \frac{\partial^2 v_\theta}{\partial z^2} \right] \end{aligned} \quad (1.3)$$

Here ρ is the density of the fluid, r represents those components in the radial direction, θ represents those components in the angular orientation, z represents those components in the longitudinal direction, r is the radius of the cylinder, v is the velocity component, g is the gravitational component, p is the pressure component, and μ is the viscosity of fluid. Though the equations may look unwieldy, they merely equate the stress in the fluid as the sum of viscous terms and pressure gradients. Meaning that a pressure difference will attempt to move a fluid and the fluid will attempt to resist this change. This is precisely how we have understood fluid flow to this point and we would do well to remember it going forward. Our use of equations will help us to systemize our reasoning.

To begin with the simplest case, let us imagine an arbitrary blood vessel in the body as a long, mostly straight, mostly cylindrical channel filled with blood. In such an environment, we can assume negligible velocities in the radial and angular directions, that the velocity in the longitudinal direction is a function of the radius, and that the effects of gravity are negligible. Furthermore let us assume that the bloods viscosity is constant and the flow is steady and incompressible. Substituting $v_r = v_\theta = 0$, we can rewrite equation 1.2 as

$$0 = -\frac{\partial p}{\partial z} + \frac{\mu}{r} \frac{\partial}{\partial r} \left(r \frac{\partial v_z}{\partial r} \right) \quad (1.4)$$

Rearranging, we can integrate and integrate again to obtain a unbounded solution for pressure driven cylindrical flow.

$$\frac{r}{\mu} \frac{\partial p}{\partial z} = \frac{\partial}{\partial r} \left(r \frac{dv_z}{dr} \right) \quad (1.5)$$

$$\int \frac{r}{\mu} \frac{\partial p}{\partial z} dr = \int d \left(r \frac{dv_z}{dr} \right) \quad (1.6)$$

$$\frac{r^2}{2\mu} \frac{\partial p}{\partial z} + C_1 = r \frac{dv_z}{dr} \rightarrow \frac{r}{2\mu} \frac{\partial p}{\partial z} + \frac{C_1}{r} = \frac{dv_z}{dr} \quad (1.7)$$

$$\int \left(\frac{r}{2\mu} \frac{\partial p}{\partial z} + \frac{C_1}{r} \right) dr = \int dv_z \quad (1.8)$$

$$v_z(r) = \frac{r^2}{2\mu} \frac{\partial p}{\partial z} + C_1 \ln(r) + C_2 \quad (1.9)$$

For an exact solution we must set boundary conditions. We can apply a no-slip condition at the vessel wall ($v_z(R) = 0$) and a no-shear condition at the centerline ($\frac{dv_z(0)}{dr} = 0$) to solve for the two unknown constants. Upon doing so we obtain a general velocity profile for this particular kind of pressure driven flow within a cylinder:

$$v_z(r) = \frac{R^2}{4\mu} \frac{\partial p}{\partial z} \left(\left(\frac{r}{R} \right)^2 - 1 \right) \quad (1.10)$$

From this we can easily derive the famous Hagen-Poiseulle equation by integrating the velocity profile described in 1.6 over the cross-sectional area of the vessel to obtain volumetric flow rate, Q .

$$Q = \int_0^R 2\pi r v_z(r) dr \quad (1.11)$$

$$Q = \int_0^R 2\pi r \left(\frac{R^2}{4\mu} \frac{\partial p}{\partial z} \left(\left(\frac{r}{R} \right)^2 - 1 \right) \right) dr \quad (1.12)$$

$$Q = -\frac{\pi R^4}{8\mu} \frac{\partial p}{\partial z} \quad (1.13)$$

As the partial derivative in space approaches the length of the vessel to be measure (that is, as $\partial z \rightarrow L$), the pressure gradient can be replaced by the ratio of the total pressure difference to the length of the vessel of interest, leaving us with a final formulation of

$$Q = \frac{\pi R^4 \Delta P}{8\mu L} \quad (1.14)$$

From this cursory examination we can see that in this sort of context, we can identify

three main components of resistance: the viscosity of the fluid, the length of the vessel, and the radius of the vessel. Resistance is directly proportional to the viscosity of blood and the length of the path it travels. Though there are many ways in which these two factors can change, on a day-to-day basis, the length of vessels and the blood's viscosity can be treated as constants. The third component, the radius of the vessel, has profound ramifications for the fluid flowing within it magnifying their effects by the power of four. Intuitively, this makes sense. The majority of frictional losses in a fluid occur along the boundary layer where the effects of viscosity are at their greatest. Decreasing the radius increases the ratio of this boundary layer to the rest of the cross-sectional area, subjecting the blood to more viscous forces at the expense of inertial forces.

The ratio of inertial (or momentum) forces to viscous forces was long ago popularized by Osborne Reynolds [8]. Though Reynolds was not the first to formulate this ratio or stress its importance (that honor should lie with George Stokes [9]), to this day the ratio bears his name. The Reynolds number is defined as

$$\text{Re} = \frac{\rho v L}{\mu} \quad (1.15)$$

where ρ is again the density of the fluid, v is the velocity of the fluid, D is the characteristic linear dimension (within a pipe, the hydraulic diameter is used), and μ is the viscosity of the fluid. For fluids of relatively constant densities and viscosities, only the velocity and the characteristic dimension affect the Reynolds number, indicating that either parameters increase aids in overcoming viscous losses. Throughout the body, the value of Re varies tremendously, from about 1 in small arterioles (where velocities are low and diameters are small) to about 4000 in the aorta (the largest vessel with the fastest velocities), indicating a wide variety of environments in which blood flows [10].

Another key characteristic affecting blood flow throughout the body occurs mere moments, mere centimeters, after blood is ejected from the left ventricle into the system circuit: the curvature of vessels. Though the approximation of vessels as long straight tubes serves a useful purpose for many applications, twists and turns both large and small exist in plenitude, and we must account for them.

When evaluating flow within curved vessels, we can turn to the work of William Reginald Dean who developed a parameter relating centrifugal forces to viscous forces, by combining the Reynolds number, Re , with a ratio of the inner tube radius, r , to the radius of curvature along the centerline of the vessel, R :

$$\text{De} = \text{Re} \cdot \sqrt{\frac{2r}{R}} \quad (1.16)$$

Fully developed flow through curved tubes tends to exhibit a net velocity that skews toward the outer wall of the bend. This is true for most arterial flow. From a few basic conservation laws (mass, energy, momentum), secondary flow will develop as a pair of counter-rotating vortices with flow in the middle of the vessel moving toward the outer wall. Flows with higher Dean number can separate along the inner wall of the curve. The presence of curvature in blood vessels ensures that forces from the flow (from the shear stress of blood along the wall or inertial forces from moment in the radial component of the vessel) will not be equal. The unequal distribution of forces along vessel walls requires the vasculature to compensate, often by thickening the lumen or increasing the inner diameter of the vessel. These effects and their role in shaping vessels (especially in the context of maturing physician created arteriovenous fistulas) will be discussed at greater length in Chapter 2.

Further variables such as the temperature of the environment, the non-Newtonian viscosity of blood, elastic wall boundaries capable of passive and active contraction and dilation, and secondary body forces, coupled with the already complex nature of continuum mechanics, all conspire to make biological fluid flow an incredibly hard problem to fully solve. However, each of these additional variables effects on flow for our present study can be considered negligible compared to that of time-variance in fluid flow.

Unsteady, pulsatile flow through nearly all of the cardiovascular system dominates many aspects of the system. From the varying pressure gradients directing blood flow through the heart and the constantly loaded and unloaded stresses applied to the walls of the arteries to the frequency component of cardiac output and the time given to fill the ventricles, the pulse of the system affects it profoundly. These effects can be mechanically transduced as by the endothelial cells of the vessel wall or physiologically manifest as with the filling time. The time variance of blood flows cannot be ignored.

To model such flows fully with the Navier-Stokes equations would be difficult and nearly impossible in all but a few cases (with still fewer interesting us). Long ago, John R. Womersley developed a nondimensional parameter analyzing unsteady Navier-Stokes equations with specific reference to biological flows to determine the ratio of unsteady forces to viscous forces [11], simplifying our work and work of many others. This Womersley number, α , is found by dividing the viscous forces into the transient inertial forces to arrive at

$$\alpha = L \sqrt{\frac{\omega \rho}{\mu}} \quad (1.17)$$

where r is the inner radius of the tube, ω is the angular frequency of the oscillating

Table 1.1: Diameter of blood vessels, the velocity of blood flowing within them, and their corresponding Reynolds and Womersley numbers. From this the wide range of diameters, velocities, Reynolds and Womersley numbers can be appreciated. All values were calculated assuming a heart beat of 75 beats per minute, a blood density of 1050 kg/m^3 (at body temperature), and a blood viscosity of 4 centipoise. Within the body, each of these parameters exist within a range, but for an order of magnitude assessment of the situation, this table should suffice.

Vessel	Diameter [mm]	Velocity [cm/s]	Reynolds number	Womersley
Aorta	25	40	2700	18
Arteries	4	10	105	2.89
Arterioles	0.03	5	0.4	0.022
Capillaries	0.008	0.03	-*	0.006
Venules	0.02	1	0.05	0.014
Veins	5	5	66	3.61
Vena Cava	30	15	1200	21.6

flow, ρ is the density of the fluid, and μ is the dynamic viscosity. Since the density and viscosity of blood remain relatively constant over time, we can see that the remaining two factor influencing the Womersley number around the body as the size of the vessels and the rate of the heartbeat. When the Womersley number is low, viscous forces dominate, making velocity profiles parabolic and with the centerline velocity oscillating in phase with the driving pressure gradient. Larger Womersley numbers (≥ 10), unsteady inertial forces dominate. The amplitude of motion due to the oscillations decreases with increasing Womersley number and the phase difference between the pressure gradient and flow grows. In this way, oscillating flows act much like low pass filters (consult [11] for a clear example of this).

A small table of typical flows, Reynolds numbers, and Womersley numbers throughout various major vessels of the body can be seen in Table 1.1.

One final significant factor influencing blood flow is the Windkessel effect, which describes the interaction of the propagating stroke volume and the compliant vessels through which it travels. The introduction of the bolus of fluid into the vasculature causes the pressure to rise during systole (recall Figure 1.3), forcing blood through the vessel (a longitudinal pressure gradient), but also applying an outward pressure to the walls (a radial pressure gradient). The vasculature being composed of compliant vessels allows for a volume expansion in response to this increased radial pressure, helping to compensate for the introduction of the stroke volume.³ The elastic vessels (mainly the aorta and large arteries) act as capacitors, storing blood during systole and discharging during diastole. In

³Indeed, it is by this mechanism that physicians are able to feel for a pulse at all!

this way, compliant arteries dampen fluctuations in blood pressure and assist with consistent tissue and organ perfusion throughout the cardiac cycle. As one ages, elastic vessels can become less compliant (hardening or becoming laden with atherosclerotic plaques), increasing pulse pressure, leading to hypertension, itself a precursor to and symptom of many cardiovascular disease (such as heart attack and stroke). For the work presented here, it is important to understand the role of the Windkessel effect as a compensatory mechanism (for the data presented on maturing fistula) while also acknowledging the complexity it adds to fluid modeling (fluid-structure interactions are difficult to simulate and replicate).

This cursory examination of fluid flow within the body excludes many deeply important aspects. Some will be discussed in other chapters (such as how fluid flow influences vascular shaping), other will not (as must unfortunately be true of any but the most fully fleshed out reviews). For the rest of this chapter, we will turn our attention to measuring fluid volumes and their movement within the human body. Specifically, we will consider the effective intravascular blood volume.

1.4 Renal function

As with many things in life, fluid volumes within the body must be balanced. Achieving the optimal distribution of molecules, ensuring the availability of electrolytes (especially the universally needed sodium and potassium), maintaining the global acid-base balance and, most importantly to this discussing, regulating fluid balance across the body are the chief responsibilities of renal physiology. Often categorized as the prominent component of the urinary system, this regulatory mechanism maintains the chemical and fluidic homeostasis. The renal system achieves this homeostasis through three main functions: excretion, separating organic waste from fluids within the body; elimination, removing waste products from the body into the surrounding environment; and regulation, maintaining consistent volume and solute concentrations of blood plasma. Through these means the urinary system plays an important role creating and sustaining intravascular volume.

The primary actors of the urinary system are the kidneys. The kidneys a pair of organs located on either side of the spine in the abdominal cavity in humans usually between T12 and L3 filter blood by separating water-soluble wastes from the rests of its contents. Beginning the process, blood is received by the kidneys through the renal artery originating along the lateral surface of the abdominal aorta. With a direct connection to the abdominal aorta, the kidneys receive a very large amount of the blood pumped through the system approximately 20-25% of the total cardiac output is directed into the kidneys. Such a large proportion of the total cardiac output emphasizes the significant role filtration plays

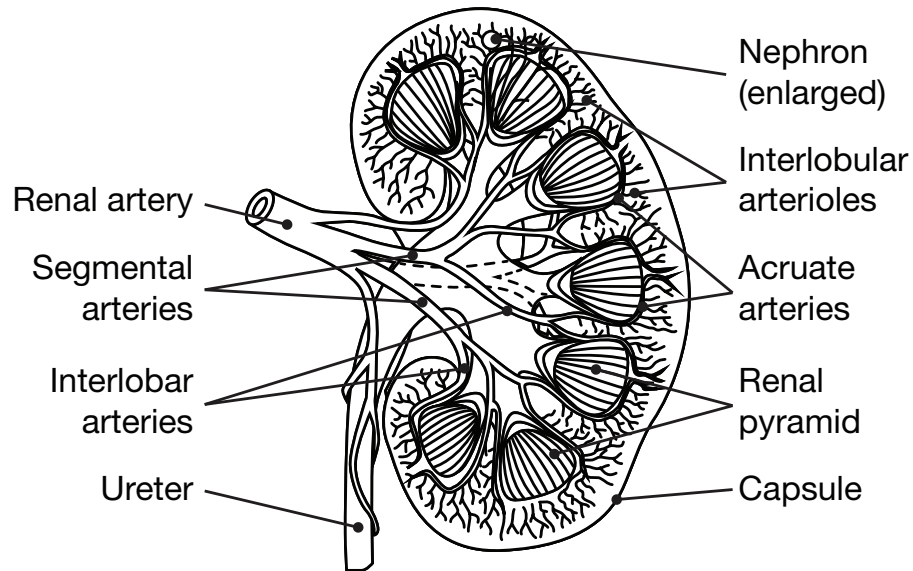


Figure 1.7: A schematic representation of the kidney. Blood flows from the renal artery into the segmental arteries to the interlobar arteries to the arcuate arteries to the interlobular arteries until finally reaching the capillary beds of individual nephrons. Once there filtration can begin with excess fluid removed, traveling from the ureter into the rest of the urinary system.

in cardiovascular regulation.

Once in the kidney, the renal artery divides several times until they are finally capillaries upon whom filtration will act. First the renal artery is divided into a series of segmental arteries that further divide into interlobar arteries. These interlobar arteries supply blood through to the arcuate arteries that lie along the boundary between the cortex and the medulla of the kidney (see Figure 1.7). The arcuate arteries divide into cortical radiate arteries (also called interlobular arteries) that then feed into afferent arterioles, which then deliver blood to the capillary beds of individual nephrons, the basic functional and structural unit of the kidney. Each nephron consists of both a renal corpuscle, a spherical structure marking the beginning of the nephron made of a glomerulus and a glomerular capsule, often called Bowmans capsule, and a renal tubule, a long slender tube-shaped structure containing the tubular fluid. Blood from the afferent arterioles is fed through to the glomerulus, a structure of a few dozen intertwining capillaries. Pressure gradients (discussed further on) force water and solutes to be filtered out of the blood and into capsular space.

The filtered product, known as filtrate, exiting the renal corpuscle is similar to protein-free blood plasma. From the capsular space, the filtrate will enter the renal tubule wherein organic materials and nutrients will be recovered, much of the water from the filtrate will




Figure 1.8: A simplified representation of renal function. Blood traveling through the afferent arteriole enters the glomerular capillaries in Bowman's capsule, where filtration occurs. Blood passing from the efferent arteriole reabsorbs and secretes fluids and materials in the peritubular capillaries, draining finally into the renal vein. All filtered fluid not reabsorbed is excreted.

be recovered, and transporting the remaining wastes further for eventual discharge. This process ensures that useful substrates present in blood plasma (such as sugars and amino acids) are retained while wastes are separated and transported further for removal. As the filtrate passes through the tubule it is transformed into tubular fluid, which gradually changes its composition until finally becoming urine. A simplified schematic of this process can be seen in Figure 1.8.

Urine, the end product of collected metabolic wastes residing in the blood, is composed of three prominent organic wastes: urea, the most abundant organic waste formed mostly from the breakdown of amino acids; creatinine, the waste product of the high energy compound, creatine phosphate, driving muscular contraction; and uric acid, the waste produced during the recycling of nitrogenous bases from RNA molecules. Each of these waste products can only be removed when dissolved in urine and thus their excretion is accompanied by unavoidable water losses. However, because the kidneys are able to produce urine ap-

proximately four times more concentrated than that of blood plasma, the fluid losses are minimized.

Here, as elsewhere, we see the presence of pressure manifest as a mode of action. Two distinct pressure gradients act in this space. The first pressure gradient, formed by the difference in the glomerular and capsule pressures, is a fluid pressure and is similar to those driving the fluids throughout the cardiovascular system. A high-pressured environment is created in the glomerulus by the arrangement of smaller efferent arteries proceeding from larger afferent arteries causing the pressure in the glomerular capillaries beds (50 mmHg) to be significantly higher than those in typical systemic capillaries (35 mmHg). Pressure in Bowmans capsule occurs from the resistance caused by the filtrate flowing through the tubule and is typically around 15 mmHg. The second pressure gradient results from an osmotic pressure difference, a phenomenon that occurs in solutions of different concentrations. When the osmotic pressure is created by differing concentrations of proteins, this is referred to as oncotic pressure. In this case, the oncotic pressures of the glomerulus and the capsule are caused by the presence and absence of blood plasma proteins, respectively. The combination of these pressures acting across the interface of the glomerulus and the capsule produce the driving force pushing the fluid into the capsular space. The rate at which fluid is filtered is known as the glomerular filtration rate (GFR) and can be represented via the Starling equation as

$$GFR = K_F[(P_G - P_B) - (\pi_G - \pi_B)] \quad (1.18)$$

where K_F is the filtration constant representing the product of hydraulic conductivity over the surface area of the acting glomerular capillaries, P_G is the hydrostatic pressure of glomerular capillary, P_B is the hydrostatic pressure of Bowmans capsule, π_G is the oncotic pressure of the glomerular capillary, and π_B is the oncotic pressure of Bowmans capsule.

Glomerular filtration is the crucial first step to all kidney function and therefore its rate must be highly controlled. Were this first step of filtration to fail many adverse consequences would follow such as wastes no longer being excreted, pH no longer remaining in balance, and a vital component of blood loss compensation would be unavailable. Three different mechanisms regulate the glomerular filtration rate: autoregulation, hormonal regulation, and autonomic regulation.

Autoregulation works at the level of the kidney and ensures that an adequate filtration rate is maintained regardless of changes in blood pressure and flow. Changes in local blood pressure trigger the arteries, arterioles, and capillaries to respond in kind to maintain a constant flow rate. Recall the fluid flow relationships described above. If a pressure drop were to be initiated, the afferent arteriole would dilate while the efferent arteriole would

constrict to counteract its effects. The converse relationship is also true. In so doing, blood flow and glomerular pressure can be kept within a consistent level regardless of systemic blood pressure.

Hormonal regulation of the glomerular filtration rate is regulated via the product of both the renin-angiotensin system (a hormone reaction occurs across the kidneys, liver, and circulatory system) and of natriuretic peptides (hormones produced by atrial and ventricular distension). A specialized set of cells in the kidneys, known as juxtaglomerular cells, synthesize the enzyme renin in response to decreased blood pressure at the glomerulus (as a result of decreased blood volume, decreased systemic pressure, or blockage of the renal artery), stimulation from sympathetic innervation, or reduced osmotic concentrations within the tubular fluid. These three phenomena are intimately related and can occur concurrently. In each case they are triggered by pressure or volumetric changes, usually within the vasculature. Renin, having been released into the bloodstream by the juxtaglomerular complex, converts an inactive protein within blood plasma, known as angiotensinogen into angiotensin I which is then converted into angiotensin II by the aptly named angiotensin-converting enzyme. This final state of the hormone, angiotensin II, affects many aspects of human physiology such as constricting the efferent arteriole in the nephron (increasing glomerular pressures and filtration rates), stimulating aldosterone release from the suprarenal glands (increasing sodium reabsorption in the tubule), and amplifies the release of antidiuretic hormone throughout the central nervous system (increasing sympathetic motor tone, initiating peripheral vasoconstriction, and mobilizing the venous reservoir). The combined effect of these responses causes systemic blood volumes and pressures to increase with the release of renin, stabilizing the glomerular filtration rate should a pressure or volume drop occur. If blood volume or pressure were to increase, however, under normal conditions the glomerular filtration rate would increase in response, promoting fluid evacuation.

Finally, autonomic innervation of the kidneys allows for sympathetic activation to override the local autoregulation control previously described. Here sympathetic stimulation has one predominant effect: it will induce vasoconstriction of the afferent arterioles, regardless of all other control. This serves to instantaneously decrease the glomerular filtration rate, allowing fluid to be retained momentarily. Sympathetic activation of this regulatory mechanism stems from acute and significant disturbances of the cardiovascular system such as a massive drop in blood pressure or a heart attack.

These three regulatory functions serve the same end: to optimize the fluid and waste removal rate to ensure a stable physiological environment.

1.5 Volume status

Extending the metaphor of the heart as a pump, the vasculature as pipes, and the blood as the fluid affected by them, we will leverage a metaphor made popular by Pereras Rapid Ultrasound in Shock evaluation protocol [12, 13] and consider another key aspect of this overall system: the tank. The tank, representing the effective intravascular volume status (here and elsewhere often shortened to simply volume status), is the measure of how much blood a person has to work with in their circulatory system at any given time. Intravascular volume is a crucial conceptual framework employed by clinicians to guide the treatment of critically and acutely ill patients. Modulation of this circulating blood volume is the foundation of many types of fluid treatment, itself a pillar of current medical practice.

One of the most important things a clinician can do for their patient is insure optimal oxygen delivery to their tissues. If we were to peak beyond the borders of Figure 1.4, we would see that the function of cardiac output itself is to drive oxygenated blood to peripheral tissues for perfusion and exchange. To sustain this feedback-control loop, water and electrolytes are necessary. People obtain both through their diet and under normal circumstances will maintain a homeostatic condition of normovolemia (or euvoolemia), sustaining a cardiac output sufficient to meet the needs of the whole body. This may, as many things can, go wrong in a number of ways.

Focusing on just those ways that in which volume status can impede proper tissue oxygenation, we can imagine that there are only two ways this can happen: a person can either have too much fluid in their blood hypervolemia or a person can have too little hypovolemia. Hypervolemia, also known as fluid overload, usually results from compromised regulatory mechanisms for electrolytes (primarily sodium) and water, as is the case in congestive heart failure, kidney failure, and liver failure (conditions responsible for millions of deaths around the globe annually [1]). Hypovolemia, also known as fluid depletion, results from large losses of volume such as vomiting, diarrhea, hemorrhage, etc. In either case, solutions exist to compensate for the fluid imbalance (sodium/fluid restriction or removal in the case of hypervolemia and infusions of replacement fluids in hypovolemia) in an attempt to optimize cardiovascular and tissue oxygenation parameters. Each solution type is predicated on altering the individual return and output curves constituting cardiac function (recall Figure 1.5).

1.6 Monitoring

Determining the presence and extent volume related conditions is crucial for diagnosing their underlying causes and to correct for them. Thus, one of the first steps in the treatment of critically and acutely ill patients is an accurate assessment of their intravascular volume.

From a clinical perspective, that is, from the perspective of someone trying to treat a patient, a more useful measure than volume status alone would be a patients volume responsiveness. Here a distinction must be made. It is not enough to merely know how full the tank is, clinicians must know how they can alter its contents (by either adding or removing fluids) and in what ways this alteration will affect the health of their patients. Therefore, in this context it is more pertinent to speak in terms of a patients potential response to a fluid than it would be to that patients overall volume. A simple enough test for responsiveness can be thought of easily: give patients some volume of a fluid and see how they respond. This is known as a fluid challenge and it gives clinicians the ability to gauge the preload reserve of patients [14], allowing them to identify those patients likely to benefit from further volume resuscitation. Such a technique is utilized in many cases, such as apparent hypovolemia, hypoperfusion, or in at-risk surgical patients [15, 16]. One must be careful in implementing and interpreting a fluid challenge, though, as research has shown that healthy volunteers can experience significant increases in stroke volume in response to virtual fluid challenges [17], suggesting that although one may respond to an initial fluid challenge, a continued fluid loading treatment may not be necessary.

Furthermore, a fluid challenge is limited to those patients whose ventricular wall could be distended from a small increase in volume. For those patients with stiffened hearts or those already hypervolemic, this type of test cannot yield any viable results of either volume status or usefully scale volume responsiveness. For those situations where a fluid challenge may not be possible or appropriate, many other techniques have been proposed and utilized throughout clinical practice.

However, several of the techniques used to predict this volume responsiveness only do so correctly about 50% of the time. Pair this with the fact that only about 50% of critically ill patients respond to volume expansion [18] and ones mind may begin to reel with the idea that clinicians need a better way of detecting volume status and volume responsiveness to treat their patients. The approach I have taken to assess volume responsiveness, one that I believe the evidence shows is better than each of those listed below, is outlined thoroughly in Chapter 3. For now, let us review the state of the field and see where each technique excels and where it may stumble.

Of the various ways clinicians have at their disposal to measure volume status and

responsiveness, multiple different kinds of classification systems that can be made. One convenient classification system that many researchers on the topic have used is dividing the field into static and dynamic measures [19, 20, 21]. The distinction here between static and dynamic is whether a signal needs to change with respect to something (time, a respiratory maneuver, etc.) to reflect volume status/responsiveness. Those that require a change are said to be dynamic, those that do not require a change, static. A further useful subdivision exists between what each of these techniques measures. In the case of volume status assessment the primary parameters of interest are pressures, volumes, and variations thereof. Pressures and volumes without variation lie within the static side of our categorization and those with variation constitute part of the dynamic side of our category.

1.6.1 The physical exam

An exception to the categories proposed above exists in that workhorse of clinical methods: the physical exam. Nearly as old as the medical field itself, the physical exam is one of the first lines of inquiry to assess many facets of patient health. The same is true when attempting to assess volume status. Studies have found the most useful parameters of the physical examination to be pulse excitation in response to postural changes (≥ 30 beats per minute) and dizziness for diagnosing hypovolemia in patients [22] and paroxysmal nocturnal dyspnea, orthopnea, and peripheral edema for diagnosing hypervolemia in patients [23]. These studies were conducted on patients with acute blood loss and heart failure, respectively, and their applicability beyond these types of patient populations is not documented. When asked to predict hemodynamic variables obtained via pulmonary artery catheterization, physicians did so correctly 24-55% of the time [24]. Such findings indicate to many that a physical examination alone is insufficient to predict the volume status or responsiveness of a patient [25].

1.6.2 Static measures of intravascular volume

To reiterate, as opposed to its more conventional meaning, a static measure of volume assessment does not mean that the signal itself does not vary with time, patient position, or stimulation. Rather, static measures are those whose values when measured are intended to provide information relevant for decision making. By way of a clarifying analogy, consider measuring the fullness of a car tire with via the pressure of air within it. Over the course of a day the overall pressure within the tire will vary as temperatures rise and fall. Even as one drives, the localized pressure distributions will change as one region of the tire experiences loading by making contact with the road. Even with all this variation, a single

static measurement of pressure can be useful to tell us whether or not we should add more air to our tires. So it goes with static measurements of volume status.

Static measures can be subdivided into either those utilizing pressures or those utilizing volumes. The principle pressure measurements that have been used for volume assessment include central venous pressure and pulmonary artery occlusion pressure. The volume measurements include right ventricular end-diastolic volume, left-ventricular end-diastolic area, inferior vena cava diameter, and bioimpedance vector analysis. Though not each of these reports a volume (one reports an area, another a diameter, the third a vector of resistance), they each represent an underlying volume of interest and are best thought of as doing so. Though there is a relationship between pressures and volumes, in the complex environments of the body (especially the broken body), they are not directly related nor can they necessarily be converted from one to the other. Hence, for examination of them we will consider their measurements separately here.

1.6.3 Pressures in general

As previously emphasize, flows stem from pressures. It makes sense, therefore, for those interested in gauging flows to turn toward pressures for guidance. For clinicians interested in where blood is, where blood is going, how it is getting there, and what it is doing once there, knowing the pressures of the environment provides insight into flow and volume related phenomenon critical for hemodynamic monitoring and assessment. Researchers of volume status and responsiveness have used and are currently using many pressures in an effort to make diagnoses and guide treatment. What follows is a short review of a few of the key pressure based techniques used to assess intravascular volume status.

1.6.3.1 Central venous pressure

Perhaps the most familiar and most hotly debated parameter for guiding fluid management is central venous pressure, the pressure of blood in the vena cava near the right atrium. Famously proposed as one of the key parameters to which therapy should be targeted [26] and a staple of international guidelines for initial resuscitation while the treating severe sepsis and septic shock [27], the reasons for central venous pressures use or non-use or its use combined with other parameters in critical care hinge on what the pressure represents physiologically, what it does to influence fluid flow, and, most importantly, what it says about the volume status of the patient.

Venous pressures arise from arterial pressures being transmitted through the capillaries beds. As the heart cannot pump out any more blood than it receives, cardiac output is

beholden to venous return. Going through the chain of causation previously outlined, we recognize that cardiac output is dependent on stroke volume which is dependent on preload which is determined by the preload, itself a factor of the filling time and venous return, which is a function between of the filling pressure and the right atrial pressure. Proponents of the metric contend that in many useful clinical settings right atrial pressure can be thought of directly through the surrogacy of central venous pressure, and therefore may be considered a reliable indicator of right ventricular preload [26, 28]. Therefore, central venous pressure is thought to give some measure of cardiac output, manifest through venous return.

Equating central venous pressure with right atrial pressure for a moment, the venous return function can be represented mathematically, first presented by Guyton [29], as

$$\text{Venous Return} = \frac{f(MCFP) \cdot f(D)}{\nu} \cdot (MCFP - RAP) \cdot C \quad (1.19)$$

where $f(MCFP)$ represents a function governing the filling pressure, $f(D)$ is a function of the dimensions of the peripheral circulatory system, ν is the viscosity of the blood, RAP is the right arterial pressure, and C is a mathematical constant for scaling the other factors. We can see from equation 1.19 that there are only a few ways of increasing venous return: we can increase mean systemic pressure, decrease right atrial pressure, or increase the difference between them. (One could also alter the arterial and/or venous resistances or the viscosity of the blood, but these are strategies not implemented to any great clinical extent current and will not concern us here.) Evidence for the existence of relationship has shown that it is an effective means of predicting the results of a fluid challenge [30].

Here confusion of the matter may begin. As measured and reported, central venous represents the pressure within intrathoracic cranial or caudal the vena cava relative to the atmosphere [28, 31]. Central venous pressure is not only not equivalent to right atrial for a great many cases, but ventricular preload is actually determined by intrathoracic transmural pressure, which is the difference between the intracardiac and extracardiac intrathoracic pressures. Transmural pressure is the more physiologically relevant parameter [31, 32] as it is actual pressure that distends the elastic structures of the heart and vessels and along whose gradients flows and forces act. Measuring the difference between the two components of transmural pressure is difficult and often not feasible in clinical practice and therefore, clinicians rely on central venous pressure as an approximation of much of the same information.

Measuring a single cannulated vessels pressure with respect to the atmosphere to obtain

Adding further to the difficulty of the matter is the fact that the information central venous pressure is meant to represent lies at the intersection of two functions of cardiac




Figure 1.9: Central venous pressured via insertion of a central venous catheter, also known as a "central line." The tip of the pressure transducer is fed through the vasculature into the superior vena cava until it is placed just outside the right atrium. A peripherally inserted catheter, or a "PICC line," can also be used to obtain central venous pressure measurements, using the peripheral vasculature to guide the pressure transducer tip to the same point in the vena cava.

performance: venous return and cardiac output (consider Figure 1.5). While the two functions are related, they can and do vary independently in a number of ways (as previously discussed).

Leaving to one side the difficulty in interpreting the bare signal of central venous pressure, a problem in consistently measuring it arises when different caregivers are asked to place the pressure transducer. Significant variability exists when different people are asked to place lines in different patients [33]. The variability in the final measurement can be greater than the magnitude of the pressure signal itself, leading to wildly different clinical decisions. As patients with the most precarious fluid balance problems are those that are most likely to have their central venous pressure monitored, it is precisely those who need an accurate measurement most that are most susceptible to errors in measurement.

Finally, there many have demonstrated that central venous pressure poorly predicts the volume status [34, 18, 35, 36], demonstrating a power of prediction right around a coins. The reasons given for central venous pressures poor predictive power can be understood from the context in which it is extracted. As I have previously mentioned, though pressure




Figure 1.10: Fed through a peripheral vein, in this case the femoral, the balloon tip of the Swan-Ganz catheter is guided through the right atrium into the right ventricle up through the pulmonary artery. When the balloon is inflated, blood flow is obstructed and the resulting pressure difference on either side, the pulmonary artery occlusion pressure, results from the pressure of the left atrium. The end-diastolic pressure in the atrium is treated as an equivalent of the pressure (and volume) in the left ventricle, and thus cardiac preload.

and resistance are two very significant contributors to blood flow (and thus blood volume shifts), many other factors including venous return to the heart, right (and left) ventricular compliance, peripheral venous tone, and patient habitus and posture all factor into the pressure measured. Furthermore, in hypovolemic patients with a functioning sympathetic response system, CVP has actually been observed to fall in response to fluid with reduced compensatory venoconstriction. It is not merely hard to account for these various confounding factors, it makes such a measurement unusable in many disease states of the critically ill such as those with pulmonary vascular disease, right and left ventricular disease and failure, and many valvular heart diseases. Given these many factors, it is possible to be fluid responsive with a high CVP and to be non-responsive with low CVP, the opposite of intuition and theory.

Nevertheless, there are those who believe that if the limitations of such a measurement are respected, some relevant information can be gleaned [28, 37]. It is still used in many clinics globally.

1.6.3.2 Pulmonary artery occlusion pressure

Pulmonary artery occlusion pressure, variously known as pulmonary artery wedge pressure, pulmonary venous wedge pressure, pulmonary capillary wedge pressure, pulmonary wedge pressure, or simply wedge pressure, provides an indirect estimate of left atrial pressure. Pulmonary artery occlusion pressure is measured by inserting a small balloon-tipped catheter (known as a pulmonary artery or Swan-Ganz catheter) into a peripheral vein (usually a jugular or femoral), feeding it through the right atrium, through the right ventricle, into the pulmonary artery. Once there, the distal end of the catheter measures pressure in the pulmonary artery (usually around 10-25 mmHg through diastole and systole). The catheter can also measure those pressures behind the pulmonary artery (from the right ventricle) and in front of the pulmonary artery (from the left atrium). Upon inflating the balloon, the pulmonary artery is obstructed, causing the pressure to drop rapidly, until it reaches a stable lower value (normally around 8-10 mmHg). All pressures behind the balloon are blocked and only those pressures from the left side of the heart are measured, specifically this occluded pressure is the left atrial pressure. During diastole when the mitral valve is open, the left atrium and the left ventricle become equally pressured. Thus, left ventricular end-diastolic pressure is equivalent to the pulmonary artery occlusion pressure. Theoretically, this left ventricular end-diastolic pressure should reflect left-ventricular end-diastolic volume (preload) and thereby was thought to be a good candidate for monitoring volume status.

Many studies have proven this untrue [4, 38, 39]. One of the assumptions about the use of pulmonary artery occlusion pressure to determine preload is that volume be proportional to pressure and that this proportionality can either be assumed or measured. Unfortunately, this pressure-volume relationship describing left ventricular compliance is often too complicated in practice for a physician to measure and too variable in context for one to simply assume. Left ventricular compliance is affected by preload, afterload, left ventricular wall thickness, elasticity, and stiffness [40, 41, 42], all of which can be radically altered in dialysis, diabetes, obesity, aging, myocardial ischemia, sepsis, etc.

1.6.3.3 Other pressures

Though still widely used, cardiac filling pressures do not appear to reliably predict fluid responsiveness [38]. Some contend that mean systemic filling pressure would be the most accurate measure of volume status [37] because of its roles in venous return. Even here, some will hedge by contending that such a measurement is not possible clinically. Others, going further, challenge this notion




Figure 1.11: Right ventricular end-diastolic volume represent the amount of blood left in the right heart at the end of diastole. The volume of the right ventricle at the end of diastole is difficult to measure with bedside techniques and is thus usually back-calculated from thermodilution based methods for determining ejection fraction.

1.6.4 Volumes in general

One of the most obvious ways one can think to measure the circulating volume in the body would be to start by measuring volumes being pumped through the heart. If one could measure the volume of blood being ejected out of the heart, the amount of blood returning to the heart, or the volume of blood/plasma through, then certainly a measure circulating volume could be gleaned. Though they specifically measure an area, a diameter, or even total body conductivity, each of the techniques that follow measure some underlying volume to assess intravascular fluid status: an area merely requires a third dimensional extension to represent volume, a diameter a long cylinder stands as a proportional measure of volume, and total body conductivity is a function of the volume of conductors.

1.6.4.1 Right ventricular end-diastolic volume

End-diastolic volume is the upper bounds of the stroke volume calculation. It is the most blood that is ever present in the ventricles and one of two parameters influencing stroke volume. Measuring end-diastolic volume should therefore be a good indicator of volume status and a key predictor of stroke volume and cardiac output. Pulmonary artery catheters outfitted with fast acting thermistors enable calculation of the right ventricular end-diastolic volume by measuring the temperature drop between two successive heart beats and corre-

lating that to the right ventricular ejection fraction. Once the ejection fraction is known, the right ventricular end-diastolic and end-systolic volumes can be calculated from the stroke volume. Though some researchers have found end-diastolic volume to accurately predict preload, cardiac indices, and more accurately guide fluid resuscitation treatments [43, 44, 45], others have not found the measure reliable [4, 38, 46]. The reason for these conflicting results is not clear, but some believe it stems from the calculation from the thermodilution-derived right ventricular ejection fraction [46].

1.6.4.2 Left ventricular end-diastolic area

Here as elsewhere, pressure-volume relationships serve as the foundation for intravascular volume. Fluids driven by venous pressure fill the heart during diastole. For a given cardiac compliance, how much fluid is driven into the ventricle should thereby be a function of the filling pressures, and how large the ventricles are should be a function of the fluid within them. In this way it is thought that if one could measure the dimensions of the ventricles, one could determine preload, and with that, gain some insight into the fluid status of the patient. Indeed, evidence exists that area of the ventricle at the end of diastole measured through transesophageal echocardiography in mechanically ventilated patients can predict volume status and responsiveness [47, 48, 49]. However, other studies have failed to replicate finding [50, 51]. One possible reason for the inconclusiveness of this technique lies in the difficulty of interpreting absolute values of the left ventricular end-diastolic area, as baseline cardiac anatomy and physiology vary widely among patients [52]. As a single static measure, ventricular dimensions cannot accurately assess where a patient lies on the Frank-Starling curve and thus have difficulty in providing metric by which decisions can be made.

1.6.4.3 Inferior vena cava diameter

There are two vessels that return deoxygenated blood from the systemic system to the heart from the upper and lower halves of the body, the superior vena cava and the inferior vena cava, respectively. Both vessels empty into the right atrium and thus the pressure in both vessels is equivalent to, or nearly equivalent to, right atrial pressure. Because the vena cavae are large, highly compliant vessels, their size can be thought to be reflective of right atrial pressure, like central venous pressure, and therefore venous return, itself an indication of volume status.

To measure the size of the vena cavae, subcostal echocardiography is employed, which in the majority of cases only allows the inferior vena cava to be visible, as seen in Figure




Figure 1.12: The inferior vena cava (IVC) as measured through subcostal echocardiography. The heart, right atrium (RA), liver, hepatic vein (HV), and common bile duct (CBD) all provide landmarks by which to make a consistent IVC diameter measurement.

1.12. Since the two vessels are roughly equivalent pressures, it does not matter which one is measured. A few researchers have demonstrated that the mean end-diastolic diameter of the inferior vena cava correlates well with mean right atrial pressure in spontaneously breathing and mechanically ventilated patients [53, 54]. Other researchers have demonstrated the ability to distinguish between hemodynamically stable and unstable patient populations based on the maximum (end-expiratory) inferior vena cava diameter [55] and could predict the recurrence of shock after fluid resuscitation better than mean arterial pressure or heart rate [56]. Unfortunately, these absolute measures of inferior vena cava diameter are essentially indirect indicators of central venous pressure and suffer from many of the same limitations.

1.6.4.4 Bioimpedance vector analysis

Related to the technique proposed in Chapter 3, bioimpedance vector analysis for analyzing fluid status has come into prominence over the past two decades because of its relative simplicity and its usefulness in key areas, such as measuring whole body fluid volumes. To use the bioimpedance vector as a metric of volume status, a functional relationship was

empirically established, by normalizing both resistance and reactance to a patient specific metric, usually by dividing these values by their height. Having performed many experiments, researchers and clinicians have been able to establish the bounds where a typical vector sought to land given a patient's hydration status. Where the tip of the normalized vector lands, can therefore indicate whether a patient has more or less fluids and/or more or less soft tissues (see Figure 3.7).

A number of studies have shown that the vector length is proportional to whole-body fluid volume with a high degree of correlation and accuracy, having been validated in many patient populations (kidney, liver, and heart failure) to assess volume status [57]. Bioimpedance vector analysis has proven especially effective evaluating the static volume status in both heart failure patients [58, 59] and for those patients undergoing dialysis.

1.6.5 Dynamic measures of intravascular volume

The body has its cycles. A great deal of attention to this point has been spent on the cardiac cycle, the beating of the heart. But there is one other cycle near this frequency range that will concern our investigations here. That is the respiratory cycle, the breathing of the lungs. Much as the purpose of the cardiovascular system is to insure that blood volumes get to where they are needed throughout the body, the purpose of the respiratory system is to insure that gases get to where they are needed throughout the body and removed from where they are not.

The most readily apparent actors in this drama are the lungs, which are able to fulfill many of their roles with the help of the pressure changes caused by a thick muscular structure known as the thoracic diaphragm. Under normal circumstance, inhalation starts with the diaphragm contracting, flattening its curved shape, increasing the volume of the thoracic cavity, inducing a negative pressure gradient and causing air to be drawn into the lungs. During exhalation, the diaphragm relaxes, returning to its dome shape, decreasing the volume of the lungs, increasing intrapulmonary pressure, and forcing air out of the lungs. This entire process takes about 2-5 seconds for average healthy adults (for an average rate of 12-30 breaths per minute).

Respiratory induced pressure changes are not merely confined to the lungs. The movement of the diaphragm and its associated pressure/volume changes affect many other aspects of the body. Relevant to the interests of the work presented here are the effects such changes have on the cardiovascular system and more specifically on the measurement of hemodynamic parameters necessary to guide fluid treatment. These respiratory induced variations have the added benefit of acting as a kind of perturbation to the system, often

in the form of cyclic loading. By dynamically exciting the system of interest, respiratory changes reveal more than a single static measurement can offer.

Dynamic respiratory responses are especially relevant in cases of mechanical ventilation, where a device assists or replaces a patient's breathing by inducing a pressure gradient within the intrathoracic cavity. Ventilation can either be through positive pressure where air (or any specified gas mixture) is pushed through the trachea into the lungs or through negative pressure most famously embodied by the iron lung, where a large vacuumed environment is formed around the outside of a patient's chest, to force air to be pulled into the lungs. If we remember that much of the effort of clinicians in emergent and intensive care settings is to facilitate oxygen delivery throughout the body, then the prominent role of mechanical ventilation (to get oxygen to and remove carbon dioxide from the lungs of patients) in these settings becomes apparent. Mechanical ventilation can thus serve the dual role of treating the patient and enabling respiratory specific measures of circulating volume.

1.6.6 Variations in general

Perturbing a system and gauging its response is one of the most established means by which to characterize said system. By inputting a known signal, one can learn how the system behaves by measuring how the signal is modulated throughout. In the dynamic measures of volume status, some portion of the cardiovascular system is subjected to a perturbation (an increased pressure, an increased volume) and an end point is measured (vital sign change, stroke volume increase). The extent to which the input independent variable (say preload) causes a change in the dependent output variable (say cardiac output) can be viewed from a volume responsiveness perspective (in this case, how one moves along the Frank-Starling cardiac output curve). Variations in signals which are already present (such as stroke volume or the diameter of the inferior vena cava) in response to a stimulus thus offer a more nuanced view of volume status than any single static measure [60].

Many researchers have uncovered a trove of good techniques for assessing volume status and responsiveness using dynamic techniques based on respiratory and selective pressure loading including stroke volume and pulse pressure variation, esophageal Doppler monitoring, echocardiographic assessment of inferior vena cava collapse, and electrical impedance based volume measurements. The effectiveness and limitations of each of these techniques are discussed below.

1.6.6.1 Stroke volume variation and pulse pressure variation

Used almost exclusively for patients receiving positive pressure ventilation, stroke volume variation and pulse pressure gauge volume status based on a simple anatomical and physiological connection between the cardiovascular and pulmonary systems. Positive pressure ventilators periodically push air into the lungs, decreasing the preload and increasing the afterload of the right ventricle. The decrease in right ventricular preload stems from the decrease in venous return caused by the increased pleural pressure during forced inspiration. The accompanying transpulmonary increase in pressure causes the afterload in the ventricle to also increase. The combined decrease in preload and increase in afterload causes the stroke volume from the ventricle to decrease, reaching a minimum at the end of inspiration. This inspiratory reduction in right ventricular stroke volume decreases left ventricular filling after a lag of about two or three beats as blood makes its trip through the pulmonary circuit. Decreased left ventricular preload reduces left ventricular stroke volume, doing so at its minimum during the expiratory period. It follows, then, that respiratory induced variations in magnitude of ventricular stroke volume indicate should indicate preload dependence [61].

Extending the logic of this assumption from stroke volume to the pulse pressure, one must contend that for a given elastic artery, the amplitude of the pulse pressure is a function of stroke volume. This is true enough for the most part and as such stroke volume variation and pulse pressure variation can be considered roughly equivalent, especially in the case of their use in assessing volume responsiveness, where they are meant to estimate where along the Frank-Starling curve a patient is (see again Figure 1.5).

Stroke volume is rarely measured directly. Instead it is usually calculated from an arterial pressure waveform, if the arterial compliance and systemic vascular resistance are known or estimated [62]. Three main systems are currently used clinically: the FloTrac system from Edwards Lifesciences, the PiCCO monitoring system from Pulsion, and LiDCO system from LiDCO. Each is invasive, requiring either a peripherally or centrally inserted arterial catheter, and each must be calibrated to make use of the pressure-volume conversion algorithms in their pulse contour analysis: the FloTrac requires patient information including demographics and physical characteristics and the PiCCO and LiDCO system require indicator dilution cardiac output measurements. The later two systems are also capable of measuring pulse pressure variation.

Though pulse pressure is routinely measured directly in clinic (via a blood pressure cuff), the methods to do so are often cumbersome and not suited for the longer term evaluation needed for variation analysis. Invasive methods for measuring arterial and venous pressures are available, indeed we have mentioned many above, but are avoided for many




Figure 1.13: A typical arterial waveform collected using photoplethysmography. The respiratory induced variation in the waveform is used to determine a subjects volume status.

reasons (to prevent infections, secondary injuries, etc.). Therefore what is often used for pulse pressure variation analysis is the waveform generated from pulse oximetric plethysmography. Pulse oximetry a technique by which to measure oxygen saturation via a photospectroscopic algorithm through a small portion of a patients body (such as a fingertip or an earlobe) allows for non-invasive measurement technique that reliably tracks the arterial pressure component [63, 64]. In addition to determining oxygen saturation, pulse oximetry gives some sense of the change in volume occurring at the region of interest, and is thus a plethysmographic technique. In the short term, the volume in the peripheral vasculature results from changes in pressure, and the vast majority of those pressure changes stem from the arterial components.⁴ In this way, the arterial waveform can be measured photoplethysmographically. Going still further, an envelope of the arterial pressure waveform can be fit to reveal a respiratory influence (see Figure 1.13). Defining the pulse pressure variability as the difference in the maximum and minimum values of the pulse pressure wave normalized to the maximum value ($(PP_{Max} - PP_{Min})/PP_{Max}$), it can now serve as a method to gauge volume status, with the mechanism of action and the logic of the approach directly analogous to the stroke volume variation case.

A meta-analysis of the literature surrounding fluid responsiveness in arterial waveform variations have shown that both stroke volume variation and pulse pressure variation to be highly predictive of fluid responsiveness (with a mean AUC of 0.84 and 0.94, respectively) [60]. Pulse pressure variabilitys higher predictive power likely spawns from its more direct measurement as stroke volume variability relies on each manufacturer-specific algorithm for determination.

However, in cases of spontaneous breathing, stroke volume variation as a method of assessing volume status falters. This is because the amplitude and shape of the intratho-

⁴In this use case pulse oximetry serves the dual role of measuring oxygen content of blood and variations in the pressure delivering it to a region.

racic pressure function differs greatly for those breathing spontaneously and those receiving positive pressure ventilation. Coupled this the unpredictable nature of spontaneous breathing and much of the benefits of stroke volume variation are thought to become untenable [65, 66]. Unfortunately for stroke volume variation as a metric, its limited reliability under spontaneous breathing has hampered its overall adoption, especially with the rise of non-sedated, non-paralyzed patients receiving respiratory support.

1.6.6.2 Changes in aortic flow velocity

To assess cardiac output, the velocity of blood flowing in the aorta can be measured via ultrasound. A small transducer is inserted through a patient's nose or mouth into the esophagus, where it is directed toward and along the descending aorta. Once there, blood velocity in the vessel can be measured via the Doppler shift. Knowing the area of the aorta, the volumetric flow can be calculated from the distribution of velocities, and thus the cardiac output determined. (Stroke volume can itself be determined by integrating the volumetric flow over time from beat to beat.) Studies in mechanically ventilated patients have demonstrated the reliability of this technique (68,69). Like stroke volume and pulse pressure variation, the variations in aortic blood flow resulting from respiratory induced pressure changes have been shown to be strongly predictive of fluid responsiveness [50, 67].

Although this technique has proven useful, it has not been broadly implemented. This is likely due to many confounding factors serving as obstacles to its adoption including the steep learning curve needed to obtain and interpret the signal of interest, the need for sedated mechanically ventilated patients, the inability to obtain continuous measurements reliably, and the actual obstacle of having to deal with a probe in the patient's esophagus for the duration of the procedure. Until this conspiracy of mitigating factors is overcome, such a technique, though useful, will not enjoy extensive use.

1.6.6.3 Respiratory induced changes in inferior vena cava diameter

One of the most promising techniques to gauge fluid responsiveness has been ultrasound measurement of respiratory induced changes in the inferior vena cava diameter. The technique works because as one breathes the shifting diaphragm causes intrathoracic pressure gradients to arise. In spontaneously breathing people, the diaphragm pulls down during inspiration, creating a significant negative pressure change in the upper half of the body and a positive pressure change in the lower half of the body. This relationship is reversed in the case of positive pressure ventilation, where a positive intrathoracic pressure accompanies inspiration. As a very compliant blood vessel, the inferior vena cava is affected greatly




Figure 1.14: Respiratory induced variations in inferior vena cava diameter as imaged through subcostal echocardiography. (Top) B-mode ultrasound images of the inferior vena cava during inspiration (left), where the diameter is at a minimum, IVC_{min} and expiration (right), where the diameter is at a maximum, IVC_{max} , in a spontaneously breathing individual. Taking an A-line distal of the hepatic vein produces an M-mode image where the respiratory (and even cardiac) variations can be seen.

by this abdominal pressure changes, collapsing under negative pressures (inspiration during spontaneous breathing, expiration under mechanical ventilation) and expanding under positive pressures (expiration during spontaneous breathing, inspiration under mechanical ventilation). The extent of collapse (or expansion) is a function of the amount of blood present in the vessel, itself a function of the amount of blood returning to the heart, and thus of the fluid status itself.

Differing from the single static measure of the inferior vena cava diameter (previously discussed), the collapsibility of the inferior vena cava has been shown to reliably guide fluid therapy [68, 69, 70, 71]. The current line of training suggests that a collapse of less than 50% in response to inspiration or a sniff suggests a normovolemic or fluid overloaded state, whereas a collapse of greater than 50% indicates fluid responsiveness. However,

recent evidence suggests 40% may be the most predictive threshold [72]. In any case, such a measurement can be taken quickly and easily at the bedside, providing actionable data in the time it takes to take a breath. Because of the relative ease in taking the measurement (usually with subcostal echocardiography), its predictive power of fluid responsiveness, and the rise of point-of-care ultrasound system, the caval index as a guide for fluid therapy has become increasingly used clinically.

Where the technique falls short is its reliance on a clinicians skill and a patients habitus. Little training and large patients make for poor results. A further limitation of the method stems from the discretization of information, where the volume status is measured only in spot checks, failing to identify trends.

1.6.6.4 Passive leg raise induced variations

With one of the major limitations to stroke volume and pulse pressure variation as dynamic measurements for volume status functionally requiring mechanically ventilated and sedated (often paralyzed) patients, clinicians and researchers have looked to other means to dynamically test for fluid responsiveness. One method whose popularity in the research community is currently growing is the passive leg raise where a patients legs are lifted above the level of their heart. Lifting the legs of a patient lying horizontally supine will cause a standing volume of fluid to be differentially subjected to gravitational forces than the rest of the body, generating a hydrostatic pressure (in simple cases found by the product of gravity, fluid density, and fluid height). Once manifest, the gravitational transfer of blood proceeds from the low pressure environments of the raised lower limbs (from the venous reservoir) to the central circulatory compartments (in the intrathoracic cavity). The raised limb increases the mean circulatory pressure, the driving pressure for venous return, and will tend to increase right cardiac preload. This type of maneuver has been referred to as an auto-fluid challenge since it performs the function of a fluid challenge (a transient increase in preload leading to an increase in cardiac output via the Frank-Starling mechanism) without the potential drawbacks of introducing a fluid into the patients system. By dynamically loading the heart in this way, various hemodynamic parameters including pulmonary artery occlusion pressure and left ventricular end-diastolic area can be altered through a passive leg raise.

By itself, a passive leg raise does nothing. It must be tied to some measure of fluid responsiveness. The end measures of choice are usually cardiac output or some sort of arterial pulse pressure monitoring. Computing the differences between baseline operation and the pressure loaded leg-raised condition of these metrics can give a sense of the fluid status of the patient. Measures tied to passive leg raises have proven to reliably predict fluid




Figure 1.15: The procedure for performing a passive leg raise. First a subject is placed in a semi-recumbent position (at 45). After a short period of time, the patient is laid supine and their legs are raised 45 of the horizon. In this way blood from the legs is pushed to the heart, performing a non-invasive fluid challenge.

responsiveness [73, 74, 75, 76, 77], with one meta-analysis suggesting the area under the receiver operating characteristics curve equal to 0.95 [78] regardless of ventilation mode or cardiac rhythm. As such, the passive leg raise as a perturbation to test the cardiac system of patients has grown as a simple yet effective means of measuring fluid responsiveness.

But not without a few major caveats. The first and most obvious caveat is that patients must have sufficient lower limbs to perform the maneuver. In cases of severe trauma or amputation, lower limbs may not be available for lifting. To my knowledge no study exists reporting the efficacy of this technique in such populations. One must not only have sufficient lower limbs, but the patient must be able to receive rapid postural changes without harm. This excludes patients with spinal injuries and the obese.

One major caveat not generally acknowledged by the literature is the difficulty in performing such a maneuver in daily practice. To be most effective, with the patient lying supine the upper half of a patient must first be raised to a 45 semirecumbent position while the lower half is kept horizontal. From this position, the upper half of the patient is lowered to the horizontal and the lower limbs are elevated 45. The upper half of the patient must be raised before the legs in order for the passive leg raise maneuver to include blood volume from the abdominal compartment to ensure a significant elevation in cardiac preload. Starting the procedure with a patient already in the horizontal position does not effectively load the heart [79], and reduces the effectiveness of the technique. In many cases, therefore, the procedure would begin with a patient resting horizontally, the upper half of the patient would be raised, the patient would sit for at least a minute, then the upper half would be lowered and the lower half raised, and from this the volume responsiveness measure of choice could be measured (but only for about a minute [76]).

Automatic beds and wedge pillows may help this process, but more often than not a caregiver of some sort would be required to lift the legs of the patient themselves, possibly

with additional assistance if the patient is particularly large (a growing concern among many). Here inconsistencies can rise such as in the pre-procedure of first sitting the patient up before laying them down or ensuring that the legs are lifted to an appropriate height or angle.

As a way to dynamically test fluid responsiveness non-invasively, the passive leg raise is an effective technique. Its popularity among researchers and a subset of clinicians has risen with good reason. However, the practical limitations of the clinic stand in the path of its full adoption.

1.6.6.5 Bioreactance and NICOM

The last of the techniques we will concern ourselves with here is the one perhaps most related to the technique proposed in Chapter 3. Bioreactance is the name of the parameter used in non-invasive cardiac monitoring (NICOM) systems utilizing the same passive electrical properties described in the Bioimpedance Vector Analysis section. Beyond being an initiative to rebrand the term bioimpedance, bioreactance as a metric distinguishes itself from typical bioimpedance measurements by being a measure of frequency modulation rather than amplitude modulation. That is, in traditional bioimpedance systems, the magnitude of the output signal with respect to the input signal is of primary concern (as indeed it is for impedimetric intravascular volume evaluation technique developed in Chapter 3). Amplitude variations stem directly from the amount of conductor being measured. Thus, given a constant control volume, the amplitude of a bioimpedance signal will only change if the conductivity of the volume changes. Bioreactance, on the other hand, is the difference in the frequency of the output signal compared to the input signal. These frequency variations stem from phase shifts from the input signal within the system and can be caused by the movement of a conductor or the increased presence of time-dependent elements, such as cellular membranes. Because the phenomenon we are dealing with here is a phasor (that is, a phase vector), the actual physical distinction between the two methods is not as great as some of the literature may suggest.

Bioreactance, as it is currently employed, is used to assess cardiac output by applying at least four electrodes to the chest and continuously measuring the thoracic impedance. Phase shifts in the impedance signal arise from large bolus of blood ejected into the aorta. The combined effects of introducing a large volume of electrically conductive material with a significant reactive component (blood) introduce time delays in the signal traversing the chest, shifting the phase, and giving rise to bioreactance.

There is evidence to suggest that the phase shifts of the bioreactance signal are well correlated with stroke volume [80, 81] and therefore holds promise for the kind of non-

invasive cardiac monitoring it underlies. In fact, several studies have shown some effectiveness in accurately tracking cardiac output and volume responsiveness [82, 83, 84, 85]. There is also evidence to the contrary [86]. Furthermore, one of the shortcomings of the technique, especially in regards to fluid responsiveness, is its requirement of including a fluid challenge of some sort to make its measurement. Most studies couple a passive leg raise (previously discussed) with a change in cardiac output to determine the volume status and responsiveness of a patient. As such it suffers from many of the same drawbacks as a traditional passive leg raise.

1.7 A summary

The function of the cardiovascular system is of great importance to a great number of people. Indeed, it quite literally lies at the heart of many matters. The system itself is complex: a four chambered, time-, volume-, and pressure-dependent pump serves two different circuits through one complete route around the body via vessels whose size, shape, elasticity, tortuosity, and orientation varies wildly with anatomy and physiology. Within this system is a complex biological fluid whose behavior is non-Newtonian and whose total viable amount within the body, regulated in part by renal function, is woefully measured to this point by a myriad of techniques with mixed accuracy and precision.

The goal of the research presented here was to demonstrate the effectiveness of a new technique for volume evaluation. Building the foundation of previous monitoring modalities and extending the logic of the physiological system under investigation, this volume assessment method uses patient specific changes in limb impedance caused by respiratory induced variations of blood volume as a surrogate for venous return. Such a technique is non-invasive, continuous, and relies on the actual volume content in the circulatory system for its measurement, rather than any indirect parameter. In this way, it is hoped that many of the sins of the past do not come to haunt this approach. We will see.

Though one can be optimistic, we must note that very little data exists to show that any of these monitoring techniques have improved patient outcomes. In reality, the most accurate measurement in the world alone will not help a patient. The measurements must be accurate, yes, but they must also be indicative of a state that can be treated, they must be interpreted as such by a caregiver, and a clinician must act in a manner that is not only correct, but done so at a rate and to a such a degree that positively affects the patient. Many things can and do go wrong along this chain from measurement to treatment, forfeiting the health and lives of hundreds, thousands of people every day. As we proceed through this dissertation, we ought to humble ourselves at the prospect of solving such an immense

problem. The author most of all. But we would also do well to remind ourselves that these problems matter to millions and millions around the globe. And we need solutions.

Here is mine.

CHAPTER 2

Ultrasound monitoring

But some strike on objects and bounce back, coming to our ears with the imitation of words.

Lucretius, *On the Nature of Things*

2.1 Ultrasound imaging in general

To truly see within a patient is the goal of every doctor. To see what ails and what heals a person (and how the ail and how they heal) is the medical profession so defined. From the symptomology first introduced by Hippocrates to the latest in nuclear medicine, many techniques have been developed to understand the internal functions and structures of human beings. Of particular relevance to the work presented here is ultrasound imaging, a method by which high frequency sound waves are used to see inside the body. Medical ultrasound machines inject acoustic waves well above the level of human hearing (hence, ultrasound) into a person's body via a transmitter and measure the echoes that bounce off the internal structures with a receiver. In most modern cases, transmission and reception occurs at the same probe, also known as a transducer. The generation, injection, reception, reconstruction, and processing of these sound waves are the types of problems sonographers, sonologists, and ultrasound engineers seek to solve as best they can.

Just as there is no one, right way to see the surface of a person, there are many ways and means of using ultrasound to view inside a subject. The general procedure used in most modern ultrasound machines, however, is approximately the same, with only subtle details changed to emphasize or minimize certain attributes of the structures imaged. First, the frequency, shape, and focus of a collection of waves to be used for imaging are determined by a central controller. This information is then relayed to a transmit beamforming controller that then sends appropriate voltages through a switching network until finally those

voltages excite the piezoelectrical elements of a transducer probe to create a series of pressure waves to be sent into the material/subject of interest. The injected sound waves then echo off internal structures with some portion of them returning back to the transducer array. The reflected pressure waves interact with the piezoelectric elements, causing them to produce a voltage potential. Depending where in the medium the waveform was reflected from and the intensity of the reflection, there will be a series of time delays and amplitude changes in the pressure waves received by the transducer. The received signals are then passed through another beamformer and summed together. To present the results in a visually pleasing manner, the envelope of signals is then detected or estimated to translate the reflected signal into an intensity signal, which can then be further processed and compressed to produce images for interpretation.

The many ways of collecting ultrasound also elicit many ways of displaying the ultrasound signals. Perhaps the most common is as a 2D image representing the cross-section mapped out by the probe as seen in Figure 2.2. This is known as a B-mode (“brightness mode”) image and it allows clinicians to quickly and non-invasively view interior structures within a subject. However, this is but one way of doing so. Of the many types of ultrasound images, but we will here concern ourselves with just four modes: A-mode, M-mode, B-mode, and Doppler mode. A-mode (“amplitude mode”) is the simplest form of an ultrasound image. A single transducer introduces sound waves into a subject and measures the echoes. In A-mode, the intensity of these echoes is plotted against the time they take to return to the transducer, an indication of the depths to which they travelled (see Figure 2.3 for an example). M-mode (“motion mode”), extends the function of A-mode by plotting many intensity vs. depth lines in succession (as seen in Figure 2.4). In so doing, events that vary with time can be tracked more directly. B-mode (as seen in Figure 2.2) parses the collected series of delays from sound waves received across multiple piezoelectric elements and transforms them into a two dimensional picture. These pictures can be played in sequence to demonstrate dynamically varying structures. Doppler imaging makes use of the Doppler effect—perceived changes in frequency due to the motion of an object—to detect and highlight motion. In clinical practice, this is most typically used to visualize blood flow within the vasculature (as is the case in Figure 2.5).

From Figures 2.2, 2.3, 2.4, and 2.5 one can begin to anticipate some of the difficulties inherent in the processing and interpretation of ultrasound signals.¹ A-mode imaging, while the closest to the real form of the acoustic data, is difficult to decipher and without many simplifying assumptions provides no anatomical or physiological data on which to operate. M-mode imaging mitigates some of the problems associated with the A-mode by

¹They are in a very real sense noise.

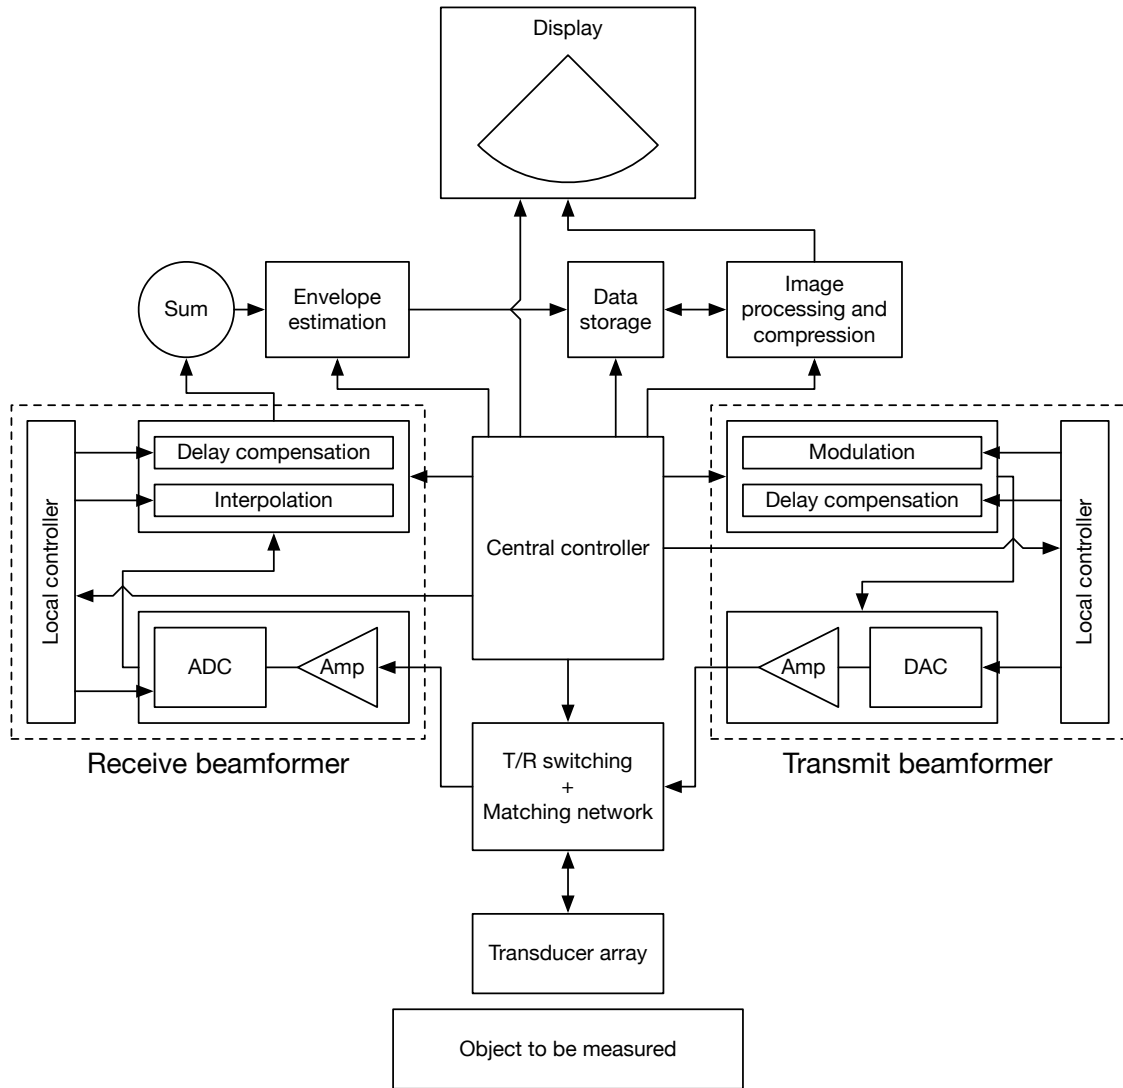


Figure 2.1: A schematic representation of a typical clinical ultrasound imaging system. The signal for a sound wave to be injected into an object through a transducer array is produced by a central controller, modulated, converted to an analog signal (through a digital-to-analog converter (DAC)), and amplified through a transmit beam former. The echoes from the input waves are received by the transducer, amplified, and converted into a digital signal (through an analog-to-digital converter (ADC)) via a receiver beam former. That digital signal is then used to estimate the envelope of the signal for further image processing and display.




Figure 2.2: Brightness mode (B-mode) ultrasound example of a post-arterial anastomotic segment of an arteriovenous fistula connected to a cephalic vein. The areas of darkness are caused by an absence of specular reflection typical of blood. The single straight line down the center of the vessel, known as an A-line, was placed by a researcher there to elucidate how A-mode and M-mode images are found (see Figures 2.3 and 2.4).




Figure 2.3: An example of amplitude mode (A-mode) ultrasound. The amplitude of the signals reflected from the medium are plotted against the time at which the signal was received. These amplitudes are taken from the A-line of Figure 2.2. Note that because of ultrasounds high frequency content (in this case the transducer used had was 7.5 MHz), the time over which these signals are collected is often less than a millisecond (and in this case is approximately $35 \mu s$).




Figure 2.4: A motion mode (M-mode) ultrasound image taken from the A-line seen in Figure 2.2. Here the slight pulsatility of the vessel can be observed in a manner that would not be obvious from a single B-mode image (as seen in Figure 2.2) or A-mode plot (as seen in Figure 2.3). Another example (in a more dynamically varying vessel) can be seen in Figure 1.14.




Figure 2.5: (Top) Doppler mode image of the arteriovenous fistula seen in Figure 2.2. A color gradient (not shown) specifies the velocity measured by the dephased acoustic waveform. (Bottom) A velocity profile can be reconstructed from the doppler image and plotted against time. Clinicians will often integrate a velocity profile found from a doppler mode ultrasound to calculate the volumetric flow.

using the signal history over time to build out a fuller picture. However, if one is not concerned with motion, M-mode is of little service. Further, there are upper and lower limits to the types of motions that can be detected: very slow and subtle motions (as those of Figure 2.4 might be to the untrained eye) and very large and rapid motions both pose problems to the clinician interested in using M-mode in certain areas. Similar critiques can be made of doppler mode ultrasound: very large and very small magnitudes of velocity measured through phase shifts can be lost or aliased. Moreover, doppler ultrasound, relying on the frequency shift of the reflected beams, can only measure motion parallel to the transducer itself. Thus, doppler sonography suffers from an angle dependence in its implementation, itself a function of the experience and intuition of the user.

B-mode ultrasound is the closest to what we traditionally think of as an image. It is a two dimensional representation of a cross-section of anatomy as viewed from a probe. Once the sound waves are processed and compressed, pixels populate a single plane to signify the intensity of the the reflected waveforms at different depths and locations within the field. Many B-mode images sequenced over time yield a video which clinicians may then play back, looking for certain features, behaviors, or attributes such as blood where it ought not to be, abnormal objects (tumors, scars, foreign bodies), or vessels that collapse too little or too much (as it is utilized for this research). Given the prevalence of B-mode ultrasound in clinic and my desire to create a clinical useful tool, the work that follows will primarily concern itself with the processing of this mode.

2.2 Theoretical foundations for ultrasound imaging

Sound describes the phenomenon of mechanical vibrations transferred through a medium through the combined effects of displacement and pressure.

2.2.1 The wave equation and its solutions

Let us begin by imagining a small volume, V , with dimensions x , y , and z , as seen in Figure 2.6. As a sound wave propagates through the medium of which the volume, V , is a part, pressures will be applied to each of the surfaces. The application of these pressures deform the volume, and this resulting volume deformation is a function of the elasticity of the material. Neglecting cases of anisotropy for a moment, we can treat the volume as consisting of a single bulk modulus, defined as the ratio of the pressure difference to the




Figure 2.6: A small volume of matter (left). When a pressure is applied to the y-z surface, the x dimension is increased by an amount Δx (center), the left most surface of the element is moved by a distance u and the right most surface is moved by a distance $u + \Delta u$ (right).

relative change in volume that results. Mathematically this can be expressed as

$$K_B = \frac{p_1 - p_2}{\frac{V_1 - V_2}{V_1}} \quad (2.1)$$

where K_B is the bulk modulus, p_1 is the original pressure, p_2 is the new pressure, V_1 is the original volume, and V_2 is the new volume. To simplify matters at this point, let us just consider the single force, F , applied to a single surface (in this case, we will use the y-z plane). As a sound wave propagates through the volume, the length of the volume along the x-direction extends from x to $x + \Delta x$. The applied force causes the surface on which it is applied to be displaced by a distance u from its initial starting point. The elasticity of the volume along the x-direction also causes an incremental displacement, Δu , to happen along its length. Assuming no energy loss, the energy inserted into the volume by the force will thus equal the potential energy stored in the volume through its elastic behavior during maximum displacement.²

Considering the case of maximum displacement³, the energy balance of the volume is thus

$$F \Delta x = K_B A \Delta u \quad (2.2)$$

where A is the area across which the force, F , is applied. Solving for force we can rewrite equation 2.2 to become

$$F = K_B A \frac{\Delta u}{\Delta x} \quad (2.3)$$

²Since we are dealing with oscillatory behavior, at all other points the energy would be a combination of potential and kinetic energies.

³Though we are considering the one dimensional case for ease of conveying the relationships, the logic of the solution that follows here can be extended into the higher dimensions.

Differentiating with respect to x and assuming a sufficiently linear system, we see that

$$\frac{\Delta F}{\Delta x} = \frac{\partial F}{\partial x} = K_B A \frac{\partial^2 u}{\partial x^2} \quad (2.4)$$

From this we have derived a relationship between the input of force on an element and its geometric extension by way of its mechanical properties. This, however, is only half of the story. Recall that a sound wave moves in space and is therefore not only a phenomenon altering the dimensions of a medium, but also transferring mass. To model this mass transport, we must turn to the ever handy second law of motion.

Newton's second law states that force is equal to the change in momentum with respect to time. That is

$$F = \frac{d(\mathbf{p})}{dt} = \frac{d(m\mathbf{v})}{dt} \quad (2.5)$$

where \mathbf{p} is momentum, m is mass, and \mathbf{v} is the velocity of the mass. In the case outlined here, the differential force, ΔF , accelerates the mass $\rho A \Delta x$ a distance u , where ρ is the density of volume. Thus

$$\Delta F = \rho A \Delta x \frac{\partial^2 u}{\partial t^2} \quad (2.6)$$

Solving for ΔF in equations 2.4 and 2.6 and combining yields

$$K_B A \Delta x \frac{\partial^2 u}{\partial x^2} = \rho A \Delta x \frac{\partial^2 u}{\partial t^2} \quad (2.7)$$

which can be further simplified to

$$\frac{\partial^2 u}{\partial x^2} = \frac{\rho}{K_B} \frac{\partial^2 u}{\partial t^2} \quad (2.8)$$

Here we can define a new parameter, c , as

$$c = \sqrt{\frac{K_B}{\rho}} \quad (2.9)$$

representing the speed of the sound wave.⁴ With this parameter, the final form of the wave equation can be written as

$$\frac{\partial^2 u}{\partial x^2} = \frac{1}{c^2} \frac{\partial^2 u}{\partial t^2} \quad (2.10)$$

It should be noted that this equation that though we have derived this equation for a lon-

⁴Such a relationship between elasticity and density can help one's intuition about the speed of sound within different materials. Materials with low bulk moduli (such as rubber) tend to have lower speed of sound than those with high moduli (such as steel).

itudinal sound wave — a wave where the force displacing the matters travels along the same path as the wave itself — it remains valid for transverse waves — waves where the direction of motion and the force are perpendicular.

There are a number of different ways to solve equation 2.2. For our purposes here we will solve the equations for pressure since it is ultimately a pressure wave that is generated and sensed by the piezoelectric elements of commercially available biomedical ultrasound machines. To begin, we will state that the displacement of the volume, V , is proportional to the force, itself a function of the applied pressure such that $u \propto p \rightarrow u = ap$, where a is a constant of proportionality relating displacement to pressure. Substituting this relationship into the wave equation yields

$$\frac{\partial^2 ap}{\partial x^2} = \frac{1}{c^2} \frac{\partial^2 ap}{\partial t^2} \quad (2.11)$$

Assuming that a remains constant over the displacements and pressures we are applying it, the above equation simplifies⁵ to

$$\frac{\partial^2 p}{\partial x^2} = \frac{1}{c^2} \frac{\partial^2 p}{\partial t^2} \quad (2.12)$$

A generalized solution to this equation is

$$p = P_0 e^{j(\beta x - \omega t)} \quad (2.13)$$

where P_0 is the initial pressure, j is the imaginary number representing $\sqrt{-1}$, β is a phase constant, x is the distance traveled, ω is the angular frequency of the sound wave, and t is the time over which the wave acts. The parameters β and ω are related to the speed of sound, c , by

$$c = \frac{\omega}{\beta} \quad (2.14)$$

itself a physical property of the medium dependent on the bulk modulus and density (as noted above).

A notable specific instance of the general solution is the expression for a forward traveling wave, of the form

$$p = P_0 \cos(\beta x - \omega t) \quad (2.15)$$

Because of the role the phase terms ωt play in this equation, to describe a backward traveling wave, one merely needs to change the $-$ sign of equation 2.15 to a $+$ sign.

⁵A further simplification that takes into consideration more dimensions is of the form

$$\nabla^2 p = \frac{1}{c^2} \frac{\partial^2 p}{\partial t^2}$$

Combining the terms from equations 2.13 and 2.14 and accounting for frequency dependent pressure variations yields another useful expression for planar waves:

$$p(x, t) = \int P(f)e^{j2\pi f(t-x/c)} df \quad (2.16)$$

Here the pressure, $P(f)$, is frequency dependent, the frequency, f , is not angular, and a depth dependent time delay can be seen, $t - x/c$. At a point $x = 0$, equation 2.16 becomes the ever familiar inverse Fourier transform

$$p(x = 0, t) = p(t) = \int P(f)e^{j2\pi ft} df = \mathcal{F}^{-1}\{P(f)\} \quad (2.17)$$

This form of the equation will become particularly useful later on.

2.2.2 Acoustic impedance and wave reflections

Recall that in most cases for ultrasound imaging to be useful clinically, some portion of the sound waves transmitted into the body must return to the transducer — there must be echoes. It is the intensity and delay of these reflections that make the visual representation of ultrasound data possible and meaningful. For a wave to be reflected back, be it a wave of light, a wave of pressure, etc., there must be some incongruity within the medium of transmission. If the medium through which a wave traveled was uniform and homogenous, nothing would alter the wave's path. For a pressure wave this incongruity arises from the coupling of two or more mediums with different characteristic acoustic impedances.

As a pressure, p , travels through a medium, it displaces particles at a velocity, v . The ratio of this pressure and ratio defines the characteristic impedance, Z_0 of a material.

$$Z_0 = \frac{-p}{v} \quad (2.18)$$

(The negative sign stems from our convention of defining a positive pressure gradient causing a velocity in the negative x-direction.) This ratio has an obvious analog to electrical impedance, which is defined as the ratio of voltage to current.

To derive the characteristic acoustic impedance of a material, let us turn to the forward-traveling wave solution we found in equation 2.15. Again, utilizing the premise that our system is sufficient linear and that displacement is again proportional to pressure, the equation can be rewritten as

$$u = U_0 \cos(\beta x - \omega t) \quad (2.19)$$

Taking the derivative with respect to time to find the velocity yields

$$\frac{du}{dt} = v = \omega U_0 \sin(\beta x - \omega t) \quad (2.20)$$

giving us the denominator of the ratio determining acoustic impedance.

Though we have an expression for pressure, its formulation with respect to the velocity equation we have derived above is not overly helpful. To help this situation let us recall that pressure is merely force acting over an area. One of our first formulations of this model (equation 2.3) will help us here. Taking the differential form and rearranging we will see that

$$\frac{F}{A} = p = K_B \frac{du}{dx} \quad (2.21)$$

where K_B is still the bulk modulus. Therefore, pressure is equal to the product of the bulk modulus and the rate of displacement with respect to the distance along the wave. By taking the derivative of 2.19 with respect to x , as in

$$\frac{du}{dx} = -\beta U_0 \sin(\beta x - \omega t) \quad (2.22)$$

and multiplying it by the bulk modulus, we would arrive at an equation for pressure that is more intuitive to work with in this context:

$$p = -\beta K_B U_0 \sin(\beta x - \omega t) \quad (2.23)$$

In this way we can represent the acoustic impedance as

$$Z_0 = \frac{p}{v} = \frac{-\beta K_B U_0 \sin(\beta x - \omega t)}{\omega U_0 \sin(\beta x - \omega t)} \rightarrow Z_0 = \frac{-\beta K_B}{\omega} \quad (2.24)$$

Substituting the results from equations 2.9 and 2.14 we are then left a simple expression for the impedance

$$Z_0 = c\rho \quad (2.25)$$

This result indicates that the characteristic impedance of a material is equal to the product of its speed of sound and the density of a material, both physical properties, both easily measured. The acoustic parameters, including the characteristic impedance, of many biological materials are listed in Table 2.1.

As seen in Figure 2.7, when a wave passes through a boundary separating two media, it can either continue on (transmit) or return back (reflect).⁶ The incident, reflected, and

⁶It may also do a number of other things, but to simplify the situation the matter here we will assume one

Table 2.1: Acoustic properties of various biological materials. Listed in this table are the density, bulk modulus, speed of sound, attenuation coefficient, and characteristic acoustic impedance for many materials that are relevant to this topic. Where applicable the mean value of a parameter for a material is accompanied by a standard deviation in parenthesis. Values are taken from Hasgall et al.

Material	ρ [kg/m ³]	K_B [GPa]	c [m/s]	α_0 [Np/m/MHz]	Z_0 [Pa-s/m ³]
Air	1.16(0.05)	0.0001	343	0.04	0.399
Blood	1050(17)	2.61	1580(11)	2.37	1660
Blood vessel wall	1102(64)	2.71	1570(6)	7.02	1730
Bone (cortical)	1908(133)	23.60	3520(420)	54.6	6710
Brain	1046(6)	2.50	1560(20)	6.80	1620
Fat	911(53)	1.89	1440(22)	4.36	1310
Heart Muscle	1081(36)	2.63	1560(15)	3.92	1690
Kidney	1066(56)	2.58	1550(18)	2.81	1660
Liver	1079(53)	2.71	1590(19)	6.92	1710
Lung	394(174)	0.36	949(342)	231	374
Mucous membrane	1102	2.89	1620(28)	120	1790
Muscle	1090(52)	2.75	1590(22)	7.11	1730
Nerve	1075(52)	2.85	1630(21)	13.2	1750
Skin	1109(14)	2.92	1620(92)	21.2	1800
Urine	1024(16)	2.42	1540(16)	0.07	1570
Water	994	2.18	1480	0.03	1470




Figure 2.7: A small volume of matter (top). When a pressure is applied to the y - z surface, the x dimension is increased by an amount Δx (middle), the left most surface of the element is moved by a distance u and the right most surface is moved by a distance $u + \Delta u$ (bottom).

transmitted waves can modeled as

$$p_1 = P_{0,1} \cos(\beta x - \omega t) \quad (2.26)$$

$$p_2 = RP_{0,1} \cos(\beta x + \omega t) \quad (2.27)$$

$$p_3 = (1 - R)P_{0,1} \cos(\beta x - \omega t) \quad (2.28)$$

Here p_1 , p_2 , and p_3 represent the pressure waves of the incident, reflected, and transmitted waves, $P_{0,1}$ is the initial pressure wave magnitude in medium 1, and R is the reflection ratio describing the proportion of the incident wave reflected back. The reflected wave travels in the direction opposite the incident wave and thus the velocities of the incident and reflected waves becomes

$$v_1 = V_{0,1} \cos(\beta x + \omega t) \quad (2.29)$$

$$v_2 = RV_{0,1} \cos(\beta x + \omega t) \quad (2.30)$$

Using the logic of equation 2.18 The wave impedance at the boundary ($x = 0$), Z_b , is the ratio of the summed pressures and summed velocities of the incident and reflected waves

$$Z_b = \frac{P_{0,1} (\cos(\omega t) + R \cos(-\omega t))}{V_{0,1} (\cos(\omega t) - R \cos(-\omega t))} \quad (2.31)$$

Since $\cos(\omega t) = \cos(-\omega t)$ and the characteristic impedance in medium 1 is $Z_1 = \frac{P_{0,1}}{V_{0,1}}$, equation 2.31 can be refined to

$$Z_b = Z_1 \frac{1 + R}{1 - R} \quad (2.32)$$

Viewing both sides infinitesimally close together, the pressures just on each side of $x = 0$ are equivalent. Furthermore, since there are no reflected waves in medium 2, the impedance at the wall is equivalent to the characteristic impedance of that medium, Z_2 , such that the equation can be reformulated as

$$\frac{Z_2}{Z_1} = \frac{1 + R}{1 - R} \quad (2.33)$$

Solving for the reflection coefficient

$$R = \frac{Z_2 - Z_1}{Z_2 + Z_1} \quad (2.34)$$

$$R = \frac{\rho_2 c_2 - \rho_1 c_1}{\rho_2 c_2 + \rho_1 c_1} \quad (2.35)$$

we see that the magnitude of a reflected sound wave is a function of two physical properties of these two events occurs.

of the tissue, namely density and the speed of sound.

2.2.3 Attenuation and enveloping

Thus far we have dealt with rather ideal media and ideal modes of transmission. In reality, the situation is far from ideal. One key concept we have heretofore neglected to mention is the loss of energy caused by a sound wave passing through a medium. In reality, energy is lost either through absorption — process in which the ultrasound energy is converted into another form of energy (such as heat, light, etc.) — or scattering — the redirection of wave energy from the path of the incident wave (through reflection, refraction, etc.) — within the medium. The collective effects of both of these energy dissipating modes is deemed the attenuation of the the wave. As can be intuited, attenuation is a multifaceted phenomenon affected by many specifics. Generally speaking, though, a medium's ultrasound attenuating behaviors are couched within a single parameter known as the attenuation coefficient, α_0 , that ties the intrinsic properties of a material to a depth and frequency dependence of the signal. As a sound wave progresses through a material, it diminishes (that is, its amplitude decreases) the further it travels.

Attenuation tends to take the form

$$\text{Attenuation} = e^{-\alpha_0 f l} \quad (2.36)$$

where, α_0 is the attenuation coefficient, f is the frequency, and l is the total path length. Going through the gamut, the pressure wave received by a pressure transducer comes from a wave reflected at a depth x , meaning that the total distance traveled in space is $2x$ and the time required to do so equal to $2x/c$. Therefore the received reflected attenuated pressure wave at the transducer becomes

$$p_x(t - 2x/c) = \int R(x) e^{-2x\alpha_0 f} P(f) e^{j2\pi f(t-2x/c)} df \quad (2.37)$$

Here we must be careful as equation 2.37 suggests something potentially disheartening. Correcting for the time delay and assuming perfect reflection makes the matter clearer

$$p_x(t) = \int e^{-2x\alpha_0 f} P(f) e^{j2\pi f t} df \quad (2.38)$$

$$= \mathcal{F}^{-1}\{e^{-2x\alpha_0 f} P(f)\} \quad (2.39)$$

$$= \mathcal{F}^{-1}\{e^{-2x\alpha_0 f}\} p(t) \quad (2.40)$$

$$= d_x(t) \cdot p(t) \quad (2.41)$$

where we have defined $d_x(t)$ as a depth dependent dispersion function. This has the unfortunate effect of suggesting that our traveling pressure waves will become ever more diffuse with distance. This does not bode particular well for signal integrity.

But there are solutions.

The current standard solution implemented in commercially available ultrasounds is to shape the pulsed ultrasound signal into a form that is less susceptible to dispersive effects. By giving the ultrasound pulse a shape, by enveloping its oscillations, the blurring due to attenuation can be mitigated, signal integrated can be maintained, and a robust system of signal detection and reconstruction can be implemented. The envelope of choice for most systems is a Gaussian pulse.

Let us introduce a simple envelope to our pressure function to see how it may combat the effects of attenuation. Working with a simple a simple Gaussian pulse envelope, our pressure function is therefore

$$P(f) = e^{-\pi((f-f_0)/W)^2} \quad (2.42)$$

where f_0 is the center frequency and W is the width of the pulse. Introducing the effects of at-

Figure 2.8: Gaussian pulse modulation of a signal. A carrier frequency signal (top) is modulated by a Gaussian envelope (middle) to produce a pulse resistant to depth dependent blurring.

tenuation, this becomes


$$P_x(f) = e^{-\pi((f-f_0)/W)^2} e^{-2x\alpha_0 f} \quad (2.43)$$

$$= e^{-2\pi f^2/W^2 + 2\pi f_0 f/W^2 - \pi f_0^2/W^2 - 2x\alpha_0 f} \quad (2.44)$$

$$= e^{-\pi \left(\left(f - \left(f_0 - \frac{\alpha_0 W^2 x}{\pi} \right) \right) / W \right)^2} e^{-2x\alpha_0 f_0 + (\alpha_0 W x)^2} \quad (2.45)$$

We can see with this solution the pulsed pressure envelope retains its width, W — integral for signal reconstruction — with the only drawback introduced being a slight shift in the center frequency by $\alpha_0 W^2 x / \pi$. It should thus not surprise us that such a solution is generally employed.

To see an example of how the carrier sound wave is modulated by the Gaussian pulse



see Figure 2.8. Another crucial feature of this pulse modulation is that it lessens the need for extraordinarily high sampling rates. Consider that even modest abdominal ultrasound probe injecting sound waves with a carrier frequency of 3 MHz must sample an rate of at least 6 MHz to even detect the signal’s oscillation, and much higher to accurately gauge the amplitude present. Subjecting the carrier signal to envelope modulation relieves this burden somewhat by making the condition for detection the proper estimation of the envelope function itself. Since pulse envelopes in most machines range anywhere from about three to six wavelengths of the carrier frequency, the sampling rate requirements for reconstruction can be greatly reduced without any loss in accuracy.

2.2.4 Envelope detection and image reconstruction

To detect the envelope of an ultrasound pulse is a relatively straightforward process. Assuming that the pressure envelope $p(t)$ is a real received signal, applying a Hilbert transform⁷ will produce an identical signal with all phase components shifted by -90. That is, the Hilbert transform acts as an ideal -90 phase shifter. Therefore, for a received signal with a Gaussian envelope, the Hilbert transform will in turn conform to a Gaussian envelope, only shifting the underlying signal. Moreover, by taking the absolute value of the Hilbert transform, the envelope signal can be entirely reconstructed. An example of the Gaussian pulse signal produced in Figure 2.8 subjected to a Hilbert transform can be seen in Figure 2.9. Scaling the amplitude of this estimated envelope to a pixel intensity from that point for image generation is one of preference, though typically some sort of logarithmic compression is used.

Though we have thus far spoken only in terms of single pulsed waves and a single reflective surfaces, it is not difficult for us to imagine that in reality many pulses are injected into materials, and that those materials are themselves composed of many possible sources of reflection. As with all waves, there is constructive and destructive interference altering

⁷The Hilbert Transform of a signal, $f(t)$, is mathematically expressed as

$$\begin{aligned} H(f)(t) &= \text{p.v.} \int_{-\infty}^{\infty} f(\tau)h(t - \tau)d\tau \\ &= \frac{1}{\pi} \text{p.v.} \int_{-\infty}^{\infty} \frac{f(\tau)}{t - \tau} d\tau \end{aligned}$$

where the function $h(t) = 1/\pi t$ is convolved with $f(t)$, p.v. is the Cauchy principal value, t is time, and τ is a delay. The principal value integral can be written explicitly for the Hilbert transform by

$$H(f)(t) = -\frac{1}{\pi} \lim_{\epsilon \rightarrow 0} \int_{\epsilon}^{\infty} \frac{f(t + \tau) - f(t - \tau)}{\tau} d\tau$$




Figure 2.9: An example of a Gaussian pulse signal subjected to a Hilbert transform. Note that what appears to be an envelope around the signal was found by taking the absolute value of the Hilbert transform, thus demonstrating its robust ability to estimate the envelope of a signal.

the received signals, though the specifics of this phenomenon will not concern us here. What all of this adds up to the establishment of a

If we consider an image space coordinate system consisting of points in (x, y, z) and point scatterers within a source system described by coordinates $(\hat{x}, \hat{y}, \hat{z})$, then the point spread function of an ultrasound system with respect to these point scatterers is

$$\text{PSF} := h_p(x, y, z : u, v, w) \quad (2.46)$$

where the location of a point scatterer is (u, v, w) in the source system. A number of simplifications can be made at this point. The first is to assume that the scattering medium can be represented by a continuous two-dimensional distribution of scatterers, represented by $s(\hat{x}, \hat{y})$ and that the corresponding reconstructed image can be represented by an analogous $I(x, y)$. The next is to assume that the point spread function is independent of the position of a scatterer within the source plane. Coupling this conclusion with the assumption that no distortion exists within the system, the image plane coordinates are linearly related to the source plane coordinates through a magnification constant, M . The resulting image is then

$$I(x, y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} s(\hat{x}, \hat{y}) h_p(x - M\hat{x}, y - M\hat{y}) d\hat{x} d\hat{y} \quad (2.47)$$

Defining slightly altered magnified variables as $\hat{x}' = M\hat{x}$ and $\hat{y}' = M\hat{y}$, equation 2.47 representing the image becomes a simple two-dimensional convolution

$$I(x, y) = \frac{1}{M^2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} s\left(\frac{\hat{x}'}{M}, \frac{\hat{y}'}{M}\right) h_p(x - \hat{x}', y - \hat{y}') d\hat{x}' d\hat{y}' \quad (2.48)$$




Figure 2.10: Ultrasound waveform envelope transformed into pixel intensity values. This data represents the first 6 mm (approximately 8 μ s) of the A-mode ultrasound waveform shown in Figure 2.3. Such a process is done for all data received by the transducer to produce a sequence of two-dimensional images. What such pixel values look like in practice may be seen in the top quarter of the M-mode image of Figure 2.4.

and as such can be more concisely written as

$$I(x, y) = \frac{1}{M^2} s(\hat{x}, \hat{y}) * * h_p(x - \hat{x}, y - \hat{y}) \quad (2.49)$$

Reconstructing the intensity for each representable point in the imaging plane, subsampling, compressing, and scaling to suit creates the images we are familiar with when discussing clinical ultrasound. The combination of many scatterers in the source plane produces the speckle patterns that so characterize ultrasound images. A source of some frustration to many, the removal of specular noise so that anatomical structures of interest can be more readily seen is an area of intense research. However, for the task that I have set for here, speckles and specular reflections are necessary to track and segment images properly.

The echoes of the input waves return to the transducer, becoming electricity once again. Once discretized, interpolated, delay compensated, and envelope detected (recall Figure 2.1), the last major thing left for an ultrasound system to do for a user is to display the information in a way that make sense. As we have seen previously in Figures 2.2, 2.3, 2.4, and 2.5, there are many such ways. Of primary importance to the image processing algorithms described below is the creation of a sequence of two-dimensional B-mode images. A single B-mode image represents the delays and acoustic impedance interfaces as a series of pixel intensities reported at a depth and lateral position, generating images like that seen in Figure 2.2, giving clinicians an intuitive way of reading the image. A basic demonstration of how pixel values may be obtained from detected envelope can be seen in Figure 2.10.

The image processing of this research works with ultrasound images after they have

become B-mode images, indeed, after they have been compressed further and stored in a standard format (the Digital Imaging and Communications in Medicine (DICOM) standard). As can be gleaned from this cursory explanation of the images we are left with, a significant amount of information has already been sapped away: analog echoes have been converted to digital electrical signals, subjected to envelope estimation, and finally pixelated for viewing on a screen. While such an end result is great for an initial inspection via ultrasound, clinical researchers wishing to utilize their machines in new and novel ways have a difficult time making images work for them. The purpose of my effort is to provide a set of post-clinical tools to work with images already collected at the bedside. The exceedingly vast majority of machines are capable of producing the sorts of images that the following analysis presents. Moreover, I present several cases demonstrating the efficacy of these tools in clinical situations, before using them in conjunction with the bioimpedance based techniques of the third chapter.

2.3 Basics of image registration

The process by which an image is transformed into another image is known as registration, and solutions for doing so are said to be solving a image registration problems [87, 88]. In practice, a sequence of images may change with time, depth, viewpoint, illumination, orientation, or any other myriad factors, with certain types of transformation distinctly measurable (for instance, illuminated scenes tend to increase regions of brightness). How to measure image transformations and what those transformations in turn tell us about what occurs in the image shall be the focus of this section.

The majority of registration methods consist of four main steps: *feature detection* wherein distinct patterns (such as lines, edges, corners, etc.) are found (either manually or automatically); *feature matching* in which detected features from a reference image are found in a secondary image; *transformation estimation*, after features are tracked and matched from one image to another their mode of transformation (the model and parameters of mapping functions) are estimated; and *image resampling* of the secondary image based on the predicted model of transformation. As the work here is primarily concerned with determining the effects of transformation on an image, the final step (image resampling) will not factor into the discussion that follows.

Specific implementation of these steps depends on the mode of image acquisition, the assumed nature of features and their transformations, and the satisfaction of the constraints of different registration techniques. A full review of image registration processes found in the medical imaging literature can be read in survey by Maintz and Viergevers [89].

The thrust of the work that follows is concerned with tracking motion in the vasculature associated with the respiratory and cardiac cycles as measured by ultrasound. I hope the reader understands my motivation, if I place an emphasis on methods pertaining to this goal at the expense of others that do not.

2.3.1 Feature detection

Just what a feature is depends almost entirely on what one wishes to do with it. It follows then that the detection of such a feature is also intimately bound up in this decision. Features can range from points (at corners, line intersections, specific places on curves, etc.) to lines (boundaries, edges, striations, etc.) to significant regions (such as lakes or forests in satellite imaging or organs and bones in x-ray computed tomography) and may be detected manually, automatically, or some algorithmic combination of the two. One may also do away with detecting features entirely and work with area-based methods that segment images into discrete units for evaluation with no regard to underlying patterns within the image (particularly useful for images with little inherent information, such as random speckles).

Though the methods used in this research rely primarily on a clinician's input to detect features of interest, a brief review of automatic detection methods ought to be instructive.⁸ A point of interest may be found through either the intersection of lines [91], corners [92], centroids of areas or volumes to be tracked, and local extrema [93]. Lines are usually detected through established edge detectors (Canny, Laplacian of Gaussian, Prewitt, Roberts, Sobel, etc.) [94, 95] relying on gradient and partial derivatives of the image to find discontinuities representing the presence of a distinct line. Regions can be found via intensity or color thresholding, pattern clustering, and closed boundary loops (found via edge detection) [96, 97].

Human beings are particularly good feature detectors from millions of years of evolution and decades of real experience benefitting those that are able to find patterns. Years of experience train clinicians to know what they are looking at and what they are looking for. Without context, experience, and a certain level of intuition, many medical images and especially those from ultrasound are difficult to decipher.⁹ As such, my approach to this aspect of the image registration process is a mixed manual and automatic approach. In one instance outlined in Section 2.3.4, a user specifies points to track and an algorithm determines whether this point satisfies certain conditions to ensure proper tracking. In another

⁸Li and Allinson's review of feature detecting algorithms [90] may be consulted should the reader find the examination here far too brief.

⁹One need only reconsult Figure 2.2 and try to figure out what the bright white line above the vessel is.

case, a region of interest may be selected by a user and automatic thresholding or segmentation will be performed (as outlined in Section 2.4.1.5). In this way, years of experience initially determine the presence and position of a feature or feature set within an image with refinement then performed to ensure that matching and tracking can be performed.

2.3.2 Feature matching

Once a feature is determined in a reference image, it must be found in a secondary image, done through a process known as feature matching. For brevity, my focus here will be exclusively on correlation-based methods and feature-based methods.

2.3.2.1 Correlation-based methods

Correlation-based methods, also called template matching, combine feature detection, matching, and even tracking into the essentially the same process. Windows of a predefined size and shape segment a part or the entirety of a reference image, then use similarity measure, namely correlation, to find the new position of window in a secondary image.

Let us quickly recall the definition of convolution between an image, $I(x, y)$, and a small subset of that image, known as a kernel, $k(x, y)$

$$I(x, y) * k(x, y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} I(u, v)k(x - u, y - v) \quad (2.50)$$

and in the discrete form for the sorts of digital images we deal with here

$$I(x, y) * k(x, y) = \sum_{i=0}^{W-1} \sum_{j=0}^{H-1} I(i, j)k(x - i, y - j) \quad (2.51)$$

where i and j count all the pixels along the width, W , and height H of the image respectively. While convolution measures how a signal would be affected by another signal (in this case), our interest lies in how similar two signals are and thus we make a slight modification to equations 2.50 and 2.51 to obtain cross-correlation

$$I(x, y) \circ k(x, y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} I(u, v)k(x + u, y + v) \quad (2.52)$$

in the continuous form and

$$I(x, y) \circ k(x, y) = \sum_{i=0}^{W-1} \sum_{j=0}^{H-1} I(i, j)k(x_i, y + j) \quad (2.53)$$

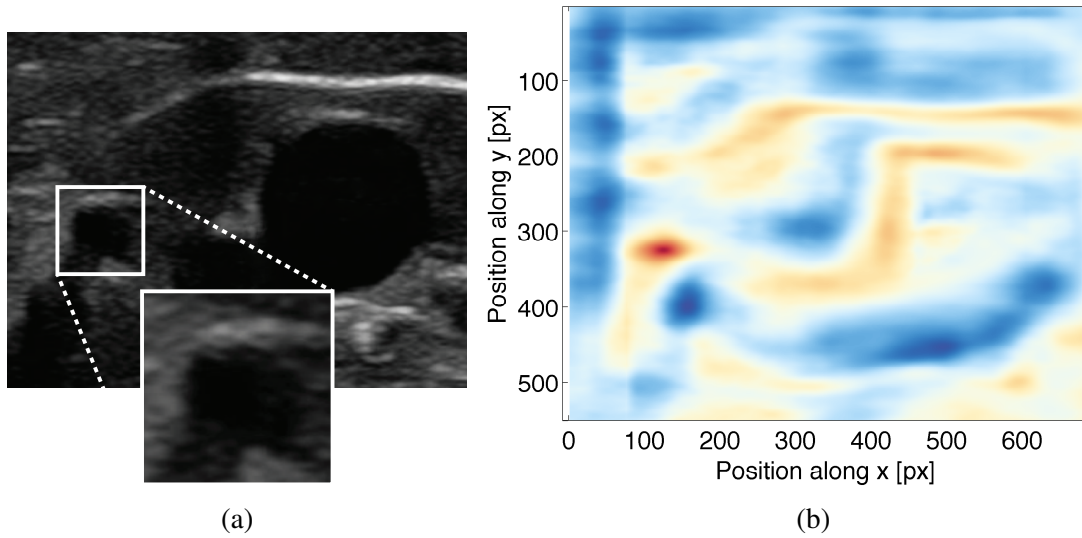


Figure 2.11: Correlation-based feature matching. (a) A subset of an image, called the kernel, is selected (in this case a portion of the post-arterial anastomotic segment from the fistula ultrasound originally shown in Figure 2.2) and its cross-correlation to every position within the search window (in this case, the entire image) can be calculated. (b) A map of the correlation coefficients at every position in the window is shown, and indeed the area of maximal correlation corresponds to the original subset image. (Red represents high values, blue represents low values.)

in the discrete form.¹⁰ Note that the only difference between the two definitions differ only in the reversal of their signs. I highlight this similarity (and difference) to emphasize the fact that many correlation techniques can make use of their relation to convolutions to increase their speed and efficiency.

A simple demonstration of correlational mapping can be seen in Figure 2.11.¹¹ In this example, the kernel window is cross-correlated across the entire image (Figure 2.11(a)), producing the correlation map of Figure 2.11(b). The point of highest correlation corresponds to the center of the kernel window chosen, showing the accuracy of the technique.

It may useful here address some of the typical concerns of using such a method. Con-

¹⁰This may also be more conveniently implemented for digital images by centering the kernel window at the point of interest:

$$I(x, y) \circ k(x, y) = \sum_{i=-w/2}^{w/2} \sum_{j=-h/2}^{h/2} I(i, j)k(x + i, y + j) \quad (2.54)$$

where w and h are the width and height of the kernel window.

¹¹Though this example uses a subset image from the reference image to demonstrate the results, the subset image would typically correlated against the secondary image.

sider again Figure 2.11(b)—one may notice a few disconcerting characteristics. First of all, though there is a single prominent peak of correlation here, this is rarely the case for images as noisy as B-mode ultrasound can be. The peak itself resides in a region of high correlation, with many of the surrounding points also very correlative with the kernel window. When small subtle changes happen quickly over such a noisy image, a blurring of correlation coefficients may happen wherein no single peak may be found. The converse is also true. Because blind correlational techniques lack any foreknowledge of a kernel’s content, they do not compensate for those windows whose content is unremarkable. That is, a kernel whose pixels describe a large smooth area without prominent details will be susceptible to incorrectly matching with other large smooth areas. Features of interest within a kernel may also appear throughout a search window, creating areas of high correlation, but of little actual relationship to the kernel itself. This may be seen in the streaks of red following the vessel’s edge, a feature that shares many similarities to the kernel (an area of bright whites next to an area of dark blacks). Precaution must therefore be taken when claiming that high correlation corresponds to a feature of interest. And finally, cross-correlation as stated here is really only suited to single straight movements of the kernel windows across the search region. Transformations more complicated than this or more myriad than uniform within the block of interest require an optimization of the kernel and search window size and shaping, ballooning the complexity rapidly.

2.3.2.2 Feature-based methods

Feature-based methods rely on two exist sets of features having already been found: one in the reference image and one in the secondary image. Such techniques attempt to find the pairwise correspondence of each point using various feature descriptors.

A straightforward way of matching a set of points in one image to a set of points in another is to assume a spatial relationship that either remains constant or subscribes to one of the transformation models outlined in the next section [98, 99]. Point- and line-matching algorithms have been developed that attempt to either minimize the error of correspondence of the displacement of points given some constraints.

A more preferable way to match features is by estimating correspondence based on the feature itself. Such techniques are said to exploit invariant descriptors. That is, the properties, descriptions, of corresponding features from a reference image to a secondary image remain the same. To be best utilized an additional set of conditions ought to be met. Namely that the features are *unique* (different features have different descriptions), *stable* (even if slightly deformed, the feature from the original image can be sensed in the secondary image), and *independent* (the elements of a feature function freely of one

another). Rarely do the stars so align that all these conditions are met. Trade-offs, assumptions, and optimizations depend on the features, the image, and the transformations expected, the specifics of which go beyond the scope of this work. (Interested readers may consult Goshtasby's *Image Registration: Principles, Tools and Methods* for a thorough review [100].) A few feature descriptions include pixel intensity, line intersections, presence within a closed-boundary region, etc.

2.3.3 Image transformation

Once a feature has been specified (or detected) in one image and matched to another, some model mapping the change in pixel data can be estimated. The goal of such a task is to find a function that corresponds to the deformation that caused the pixel data to change from one image to the next, all while cognizant of the method of image acquisition and attempting to compensate for distortions possibly introduced in the feature matching step. Much like feature matching techniques, parameters of such functions ought to be those that maximize the correspondence of the tracked features, done by minimizing the errors made of a given transformation function. Such mapping functions can either be specified in advance if an assumed geometric deformation is expected from secondary images¹² or may be systemically tested to find one that best fits the data. There are many ways to classify the types of transformation models, but for our purposes here we will consider two broad classes: global and local techniques. This distinction may also, in this context, be thought of as linear and nonlinear transformations. Though the distinction between these two types of mapping is not always clear (interested readers may consider the case of radial basis functions), for the work described here it should suffice.

2.3.3.1 Global (linear) transformations

Global techniques use all matched features to predict the model of transformation. Such methods may also be thought of as “rigid” or “linear” as all points are treated as though affected by the same mapping function. The general presentation of such a mapping function is known as an affine transform and consists of four possible combinations of transformation: scaling, translation, rotation, and shear.¹³

¹²In special cases where the geometric deformation is predetermined and well known, pre-correction and transformation may be performed.

¹³Not considered here are the effects of projective transformation as it is not relevant to the ultrasound imaging methods used.

Scaling

One of the simplest types of image transformation is scaling where the distance between points in one image are modified by some constant to yield the points in another, either uniformly or non-uniformly. For instance, if a point $\mathbf{p} = (x, y)$ in a two-dimensional¹⁴ reference image, I_1 , is related to another point $\mathbf{p}' = (x', y')$ in the secondary image, I_2 , by a constant factor, s , then we may model this as

$$\begin{aligned}x' &= sx \\ y' &= sy\end{aligned}$$

or in matrix form as

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} s & 0 \\ 0 & s \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \quad (2.55)$$

And if the scaling factors for x and y are not the same, we can modify equation 2.55 to become

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} s_x & 0 \\ 0 & s_y \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \quad (2.56)$$

where s_x and s_y are scaling factors for x and y , respectively.

Translation

If the points of an a reference image are shifted uniformly in a secondary image, then the points are said to have undergone a translation. Here the the point $\mathbf{p} = (x, y)$ (which may also be represented as $\mathbf{p} = \begin{bmatrix} x & y \end{bmatrix}^T$) are modified by a vector $\mathbf{T} = \begin{bmatrix} t_x & t_y \end{bmatrix}^T$ such that

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} t_x \\ t_y \end{bmatrix} = \begin{bmatrix} x + t_x \\ y + t_y \end{bmatrix} \quad (2.57)$$

the pixels are said to have been translated from (x, y) to (x', y') .

Rotation

Images may also be rotated about a point in their transformation. Defining the magnitude of the vector \mathbf{p} with respect to the point of rotation as $|\mathbf{p}|$ and the angle with respect to the horizontal axis as α then a rotation of an image by an angle, θ , maps the point (x, y) to

¹⁴The extension into higher dimensions follows the same logic outlined here.

(x', y') by

$$\begin{aligned}
 \begin{bmatrix} x' \\ y' \end{bmatrix} &= \begin{bmatrix} |\mathbf{p}| \cos(\alpha + \theta) \\ |\mathbf{p}| \sin(\alpha + \theta) \end{bmatrix} \\
 &= \begin{bmatrix} |\mathbf{p}|(\cos \alpha \cos \theta - \sin \alpha \sin \theta) \\ |\mathbf{p}|(\sin \alpha \cos \theta - \cos \alpha \sin \theta) \end{bmatrix} \\
 &= \begin{bmatrix} x \cos \theta - y \sin \theta \\ x \sin \theta + y \cos \theta \end{bmatrix}
 \end{aligned}$$

and thus finally

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \quad (2.58)$$

Shearing

The final linear transformation is known as shearing. Here points are displaced in a fixed direction by an amount proportional to their distance from a line parallel to their translation. Mathematically this is similar to equation 2.56 with two notable differences: the scaling is uniform and the original 0 value terms have the shear scaling constants, Sh_x and Sh_y , thus becoming

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} 1 & Sh_x \\ Sh_y & 1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \quad (2.59)$$

Shear mapping is distinct from both scaling and rotation as all lines not parallel to the axis of motion and all angles will change, though the alignment and relative distances of collinear points are preserved.

Affine transformation

Each of these types of transformation may be present for any given reference and secondary image pairing:

$$\begin{matrix} \mathbf{p}' & \text{Scaling} & \text{Translation} & \text{Rotation} & \text{Shearing} & \mathbf{p} \\ \begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} & = \begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0 & 0 & t_x \\ 0 & 0 & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & Sh_x & 0 \\ Sh_y & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} & (2.60) \end{matrix}$$

Concatenating each of these matrices (scaling, translation, rotation, shearing)¹⁵ of equation 2.60 produces a simple matrix known as an affine transformation:

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \quad (2.61)$$

The properties of the affine transform include the preservation of parallel lines, the preservation of equispaced points along lines (equally spaced points along a line in the reference image will be equally spaced along a line in the secondary image), and the preservation of the intersection of lines (also known as the incident preservation). As is apparent from equation 2.61, there are six degrees of freedom, and thus requires at least three coordinates to fully define. Global mapping schemes must therefore have at least three features matched and tracked for a model's transformation function to be fully defined.¹⁶

Typically the number of points tracked is higher than the minimum number required to determine the mapping function and as such the parameters of the function are found through some fitting function, usually by a linear regression method (a least-square fit, minimizing the sum of squared errors for the feature points). Though this has the effect of not mapping points to their counterparts directly, it is on the whole more accurate for global transformations. However, errors can accrue greatly when not all points move together linearly. In such cases, local transformation functions must be considered.

2.3.3.2 Local (nonlinear) transformations

Images that deform differently across their environment are referred to as locally deformed images. Such transformations may also be thought of as nonlinear or nonrigid [101]. These happen quite frequently in medical imaging and are thus worth our attention here. Consider the case of a viewing the cross-section of a blood vessel with an ultrasound. The vessel itself might pulse as the heart beats, but the rest of the surrounding tissue may remain still, or deform less severely. Global mapping techniques have no means by which to account for this behavior. Thus local mapping models must be used. Several techniques are available

¹⁵To be clear, the order of the transformations do matter, though I have treated them here as if they did not. For instance, rotating and then translating is not equivalent to translating and then rotating.

¹⁶A simplified transform, known as the similarity transform, consisting of scaling, translation, and rotation only, defined by

$$x' = s(x \cos \theta - y \sin \theta) + t_x \quad (2.62)$$

$$y' = s(y \cos \theta + x \sin \theta) + t_y \quad (2.63)$$

only requires two points to fully define the mapping function.

to find the parameters of such locally varying transformations, usually by modifying the fitting function of the points to a weighted least-square or weight-mean [102]. Piecewise methods are also available [103, 104] to help fit parametric mappings via interpolation.

One may also do away with parametric mapping functions entirely, as originally proposed by Bajcsy and Kovai [105]. To do so is to employ an elastic registration process where images are treated as if they were rubber sheets that are stretched and strained by many external forces counteracted by external internal stiffnesses. For elastic registration techniques, feature matching and mapping are done simultaneously. This is usually done by iteratively locating a minimum error state.¹⁷ Elastic registration techniques fall short when image deformations are very localized, requiring an approach that models the image not as a rubber sheet, but as a viscous fluid. So-called “fluid registration” processes are available [106], but are known to introduce artifactual blurring [106].

Optical flow registration techniques seek to remedy this problem by finding the relative motion between two images at every feature (usually some subset of regularly spaced pixels) and do so by assessing the apparent motion of brightness patterns in an image which are ideally the same as the motion field itself.¹⁸ If we consider a point (x, y) on a reference image with an intensity field, I , occurring at time, t , that moves a distance, $(\Delta x, \Delta y)$, over the course of some small amount of time, Δt , then we can state that

$$I(x, y, t) = I(x + \Delta x, y + \Delta y, t + \Delta t) \quad (2.64)$$

if the brightness of the feature remains constant. We may then apply a Taylor series approximation to the right hand side of the equation to get

$$I(x + \Delta x, y + \Delta y, t + \Delta t) = I(x, y, t) + \frac{\partial I}{\partial x} \Delta x + \frac{\partial I}{\partial y} \Delta y + \frac{\partial I}{\partial t} \Delta t \quad (2.65)$$

neglecting higher order terms. Substituting this result into equation 2.64 we see that

$$\frac{\partial I}{\partial x} \Delta x + \frac{\partial I}{\partial y} \Delta y + \frac{\partial I}{\partial t} \Delta t = 0 \quad (2.66)$$

which is

$$\frac{\partial I}{\partial x} V_x + \frac{\partial I}{\partial y} V_y + \frac{\partial I}{\partial t} = 0 \quad (2.67)$$

where the velocity in the x and y directions, $V_x = \frac{\Delta x}{\Delta t}$ and $V_y = \frac{\Delta y}{\Delta t}$, respectively comprise

¹⁷Such as by pyramidal segmentation discussed in the *Iterative Kanade-Lucas-Tomasi feature tracking algorithm* section.

¹⁸It is important to note that anything altering intensity (such as different lighting patterns) will produce apparent motion even though there might not be any.

the velocity vector, \mathbf{V} . This may also be rewritten as

$$\nabla I \cdot \mathbf{V} = -I_t \quad (2.68)$$

where ∇I is the image intensity gradient and I_t represents the partial derivative of intensity with respect to time.¹⁹

As may be apparent, equation 2.68 has two unknowns at a given point and thus cannot be solved directly. Optical flow algorithms must therefore apply additional constraints to determine motion. Of the many algorithms for doing so, we can focus on three broad categories: differential methods, block-based methods, and phase-based methods. A fuller review including the relative performance of such algorithms can be found in Barron et al. [107]. Only the former two aforementioned techniques will bear on the work presented here and thus the two sections that follow delve deeper into them alone.

Differential methods

Differential methods attempt to calculate velocities from the (partial) spatiotemporal derivatives of image intensities from the reference image to the secondary image. This requires, at a minimum, that the images be differentiable. To avoid aliasing or exaggerating potential problems in numerical differentiation, often some sort of temporal smoothing is applied. When large or nonlinear velocities are present, coarse-to-fine techniques for motion estimation will often be applied, where initial guesses in subsampled images are fed forward to higher resolution images.

As previously mentioned, optical flow techniques must constrain equation 2.68 to solve it. One means of doing so is to perform higher order differentiation to constrain to the two-dimensional velocity²⁰:

$$\begin{bmatrix} I_{xx}(x, y, t) & I_{yx}(x, y, t) \\ I_{xy}(x, y, t) & I_{yy}(x, y, t) \end{bmatrix} \begin{bmatrix} \mathbf{V}_1 \\ \mathbf{V}_2 \end{bmatrix} = - \begin{bmatrix} I_{tx}(x, y, t) \\ I_{ty}(x, y, t) \end{bmatrix} \quad (2.69)$$

Equation 2.69 can be used in isolation or with equation 2.68 to determine or over-determine the resulting velocities, respectively. As is a common problem for many higher order examinations, such an approach is particularly sensitive to errors stemming from numerical differentiation and are thereby often less accurate than estimates found via first-order methods [107].

¹⁹This notation for derivatives may be odd for some readers, but as it is commonly employed in the field of computer vision and image processing, I will use it throughout.

²⁰In this case we assume the conservation of $\nabla I(x, y, t)$ such that $\frac{d\nabla I(x, y, t)}{dt} = 0$.

Another more common way to constrain the velocity function is by combining local estimates with broader two-dimensional velocity fields through space and/or time [108, 109]. This may be done either through a global smoothness constraint (as is the case in a technique like the Horn and Schunck method [110]) or by fitting regions to local models of velocity (as in the Lucas and Kanade method [111]).

The Horn and Shunck method seeks to find a velocity field, $\mathbf{V}(x, y, t)$ (equivalent in this case to $(u(x, y, t), v(x, y, t))$) by postulating a global energy function, then seeking to minimize it. Here the energy, E , is defined as

$$E = \iint ((\nabla I \cdot \mathbf{V} + I_t)^2 + \lambda^2(\|\nabla u\|^2 + \|\nabla v\|^2)) dx dy \quad (2.70)$$

over the image of interest. The parameter λ is a regularization constant and larger values are thought to lead smoother flows (per Horn and Schunck [110]). This function can be solved numerically by solving a pair of Euler-Lagrange equations (for two-dimensional flows) from the calculus of variations:

$$\frac{\partial L}{\partial u} - \frac{\partial}{\partial x} \frac{\partial L}{\partial u_x} - \frac{\partial}{\partial y} \frac{\partial L}{\partial u_y} = 0 \quad (2.71)$$

$$\frac{\partial L}{\partial v} - \frac{\partial}{\partial x} \frac{\partial L}{\partial v_x} - \frac{\partial}{\partial y} \frac{\partial L}{\partial v_y} = 0 \quad (2.72)$$

where L is an integrand of the energy equation (2.70), yielding

$$I_x(I_x u + I_y v + I_t) - \lambda^2 \left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} \right) = 0 \quad (2.73)$$

$$I_y(I_x u + I_y v + I_t) - \lambda^2 \left(\frac{\partial^2 v}{\partial x^2} + \frac{\partial^2 v}{\partial y^2} \right) = 0 \quad (2.74)$$

where, again, I_x , I_y , and I_t represent the partial derivatives of image intensity with respect to the x, y, and time dimensions, respectively. Since this is meant to be a numerical derived values, the Laplacian operations $\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2}$ and $\frac{\partial^2 v}{\partial x^2} + \frac{\partial^2 v}{\partial y^2}$ may be rewritten as a weighted average based on regional values augmented by a point specific velocity:

$$\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} = \bar{u}(x, y) - u(x, y) \quad (2.75)$$

$$\frac{\partial^2 v}{\partial x^2} + \frac{\partial^2 v}{\partial y^2} = \bar{v}(x, y) - v(x, y) \quad (2.76)$$

Combining these results with 2.74 and combining terms yields equations of the form

$$(I_x + \lambda^2) u + I_x I_y v = \lambda^2 \bar{u} - I_x I_t \quad (2.77)$$

$$(I_y + \lambda^2) v + I_x I_y u = \lambda^2 \bar{v} - I_y I_t \quad (2.78)$$

which Horn and Schunck suggested be found through iteration as the solution for point specific velocity is dependent on the region in which it is placed. An iterative method can be derived from this formulation of a Jacobian nature

$$u^{k+1} = \bar{u}^k - \frac{I_x(I_x \bar{u}^k + I_y \bar{v}^k + I_t)}{\lambda^2 + I_x^2 + I_y^2} \quad (2.79)$$

$$v^{k+1} = \bar{v}^k - \frac{I_y(I_x \bar{u}^k + I_y \bar{v}^k + I_t)}{\lambda^2 + I_x^2 + I_y^2} \quad (2.80)$$

where k denotes the last calculation and $k + 1$ represents the next iteration. This approach has the advantage of handling high density flow fields, but is susceptible to highly localized variations. Presmoothing and precaution ought to accompany its use.

The Lucas and Kanade method begins by assuming features within a local region of a point of interest experience very small displacements, such that the contents of the image are approximately constant. With this constraint, equation 2.68 can be assumed to hold for a local region of pixels within a window centered at some point of interest, \mathbf{p} so that the velocity vector $\mathbf{V} = (u(x, y, t), v(x, y, t))$ is satisfied at all points (q_1, q_2, \dots, q_n) :

$$\begin{aligned} I_x(q_1)u + I_y(q_1)v &= -I_t(q_1) \\ I_x(q_2)u + I_y(q_2)v &= -I_t(q_2) \\ &\vdots \\ I_x(q_n)u + I_y(q_n)v &= -I_t(q_n) \end{aligned} \quad (2.81)$$

This may be rewritten in matrix notation as

$$A \cdot \mathbf{V} = b \quad (2.82)$$

where

$$A = \begin{bmatrix} I_x(q_1) & I_y(q_1) \\ I_x(q_2) & I_y(q_2) \\ \vdots & \vdots \\ I_x(q_n) & I_y(q_n) \end{bmatrix}$$

$$\mathbf{V} = \begin{bmatrix} u \\ v \end{bmatrix}$$

and

$$b = \begin{bmatrix} -I_t(q_1) \\ -I_t(q_2) \\ \vdots \\ -I_t(q_n) \end{bmatrix}$$

With a region of any appreciable size, such a system of equations immensely over-determines the problem. The Lucas and Kanade method resorts to a least squares fitting method to solve the velocity:

$$\mathbf{V} = (A^T A)^{-1} A^T b \quad (2.83)$$

where A^T is the transpose of matrix A . In matrix form this becomes

$$\begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} \sum_{i=1}^n I_x(q_i)^2 & \sum_{i=1}^n I_x(q_i)I_y(q_i) \\ \sum_{i=1}^n I_y(q_i)I_x(q_i) & \sum_{i=1}^n I_y(q_i)^2 \end{bmatrix}^{-1} \begin{bmatrix} -\sum_{i=1}^n I_x(q_i)I_t(q_i) \\ -\sum_{i=1}^n I_y(q_i)I_t(q_i) \end{bmatrix} \quad (2.84)$$

Such a result may be further refined by subjecting the window for observation to weighting function, w , that preferentially favors some portion of the window (for instance, the center): $\mathbf{C} = (A^T w A)^{-1} A^T w b$. Many finite impulse response filters can and have been used (such as a simple Gaussian function). Infinite impulse response filters applied recursively have also been used successfully.

Block-based methods

Block-based methods, variously called region-based [107] or block-matching [112], begin by defining velocity as a shift of a feature by a displacement vector $\mathbf{d} = (d_x, d_y)$ over the frames (and thus time) that produces the best fit between two regions. The best fit is generally defined by one of two similarity measures by either maximizing the correlation (such as by the normalized cross-correlation function) or by minimizing the differences between the two regions (as by the sum-of-squared-differences). These two similarity functions are defined as

$$NCC = \frac{\sum_{(i,j) \in W} I_1(x+i, y+j) \cdot I_2(x+d_x+i, y+d_y+j)}{\sqrt{\sum_{(i,j) \in W} I_1^2(x+i, y+j) \cdot \sum_{(i,j) \in W} I_2^2(x+d_x+i, y+d_y+j)}} \quad (2.85)$$

$$SSD = \sum_{(i,j) \in W} (I_1(x+i, y+j) - I_2(x+d_x+i, y+d_y+j))^2 \quad (2.86)$$

respectively, where i and j designate all points in the window, W , I_1 is the intensity map of the region in the reference image, and I_2 is the intensity map of the region from the secondary image. Minimizing the sum-of-squared differences essentially maximizes the integral of the product $I_1(x, y)I_2(x + d_x, y + d_y)$ over the window and can thus be viewed as a window-weighted first-order approximation of the temporal derivative, I_t .

The simplest block-matching algorithm to perform is known as a full, or exhaustive, search. It is computationally expensive as it attempts to calculate the similarity function at every possible location within a search window. The advantage of such an approach is the comfort in having tested every single possible displacement a region may have been subjected to and because of the methods relation to cross-correlation, and thus convolution, easy to implement in hardware. Obvious disadvantages include that even for very small features within small search windows calculating every possible motion takes an immense amount of computation—in fact, the number of calculations is the same regardless of the presence or extent of motion. Several faster searching techniques have been developed (such as multi-step and predictive searches) to mitigate computational complexity by implementing course-to-fine segmentation, assuming patterns of local variation, etc., doing so with varying success [107, 112].

2.3.4 Iterative Kanade-Lucas-Tomasi feature tracking algorithm

The algorithm used for this research makes use of two techniques, the Kanade-Lucas-Tomasi (KLT) feature tracker and pyramidal segmentation. The KLT feature tracker is a technique commonly used in computer vision to follow certain image features (edges, points, etc.) from one frame to the next. Traditional image tracking techniques can be computationally costly as they try to match a subset of pixels of interest (known as a kernel) from one frame to within another, larger subset of pixels (known as a search window) in a frame that follows. There are many techniques to perform this action with varying costs of computation. To mitigate these computational costs the KLT technique makes use of spatial intensity information to direct the search for the motion vector that best describes the features change in position from one frame to the next.

One of the limitations of the KLT tracker is that the motion vector must be small. That is, the KLT technique alone is not capable of measure large changes in position. To work around this problem, our algorithm employs a pyramidal segmentation technique to subsample an image reducing its resolution and increasing the spatial information inherent in each pixel to perform large motion tracking, then feeding the results of the low resolution tracked images to higher resolution images. In this way the accuracy and speed of the KLT

technique can be used to track a feature, in this case the walls of the IVC, over large distances.

To begin tracking, a pair of points on the vessels edge are selected on an initial image. The two points are tracked independently as follows. Let I_1 and I_2 be two grayscale images in sequence. $I_1(x, y)$ is the pixel intensity value of I_1 at a point, (x, y) . Let $\mathbf{u} = (u_x, u_y)$ be a point on the first image, I_1 . The goal is to find a point \mathbf{v} on I_2 where $I_1(\mathbf{u})$ and $I_2(\mathbf{v})$ are most similar, with the assumption that the feature located at $I_1(\mathbf{u})$ has moved to $I_2(\mathbf{v})$ by a motion vector $\mathbf{d} = (d_x, d_y)$. To establish the frame-to-frame motion vector, \mathbf{d} , we will seek to minimize a residual function, a sum of squared differences, defined as

$$\begin{aligned}\epsilon(\mathbf{d}) &= \epsilon(d_x, d_y) \\ &= \sum_{x=u_x-w_x}^{u_x+w_x} \sum_{y=u_y-w_y}^{u_y+w_y} (I_1(x, y) - I_2(x + d_x, y + d_y))^2\end{aligned}\quad (2.87)$$

where $\epsilon(\mathbf{d})$ is the residual error, w_x and w_y are the integration window size parameters (with a window size equal to $(2w_x + 1) \times (2w_y + 1)$), centered at point (x, y) . If we then set the derivative of the matching error with respect to the displacement vector to 0, the error minimizing displacement vector can be found.

$$\begin{aligned}\frac{\partial \epsilon(\mathbf{d})}{\partial \mathbf{d}} &= 0 \\ &= -2 \sum_{x=u_x-w_x}^{u_x+w_x} \sum_{y=u_y-w_y}^{u_y+w_y} (I_1(x, y) - I_2(x + d_x, y + d_y)) \begin{bmatrix} \frac{\partial I_2}{\partial x} & \frac{\partial I_2}{\partial y} \end{bmatrix}\end{aligned}\quad (2.88)$$

Simplifying equation 2.88 by substituting the first order Taylor series expansion of $I_2(x + d_x, y + d_y)$ gives

$$\frac{\partial \epsilon(\mathbf{d})}{\partial \mathbf{d}} \approx -2 \sum_{x=u_x-w_x}^{u_x+w_x} \sum_{y=u_y-w_y}^{u_y+w_y} \left(I_1(x, y) - I_2(x, y) - \mathbf{d} \begin{bmatrix} \frac{\partial I_2}{\partial x} & \frac{\partial I_2}{\partial y} \end{bmatrix} \right) \begin{bmatrix} \frac{\partial I_2}{\partial x} & \frac{\partial I_2}{\partial y} \end{bmatrix}\quad (2.89)$$

From here we can define the frame-to-frame difference at the point of interest as

$$\delta I(x, y) = I_1(x, y) - I_2(x, y)\quad (2.90)$$

and the image gradient as

$$\nabla I = \begin{bmatrix} I_x \\ I_y \end{bmatrix} = \begin{bmatrix} \frac{\partial I_2}{\partial x} & \frac{\partial I_2}{\partial y} \end{bmatrix}^T\quad (2.91)$$

where

$$I_x(x, y) = \frac{\partial I_1(x, y)}{\partial x} = \frac{I_1(x + 1, y) - I_1(x - 1, y)}{2} \quad (2.92)$$

$$I_y(x, y) = \frac{\partial I_1(x, y)}{\partial y} = \frac{I_1(x, y + 1) - I_1(x, y - 1)}{2} \quad (2.93)$$

Substituting into 2.89 yields

$$\frac{1}{2} \frac{\partial \epsilon(\mathbf{d})}{\partial \mathbf{d}} \approx \sum_{x=u_x-w_x}^{u_x+w_x} \sum_{y=u_y-w_y}^{u_y+w_y} (\mathbf{d} \nabla I^T - \delta I) \nabla I^T \quad (2.94)$$

Combining terms and transposing creates

$$\frac{1}{2} \left[\frac{\partial \epsilon(\mathbf{d})}{\partial \mathbf{d}} \right]^T \approx \sum_{x=u_x-w_x}^{u_x+w_x} \sum_{y=u_y-w_y}^{u_y+w_y} \left(\mathbf{d} \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix} - \begin{bmatrix} \delta I & I_x \\ \delta I & I_y \end{bmatrix} \right) \quad (2.95)$$

Defining a G matrix and b as

$$G \doteq \sum_{x=u_x-w_x}^{u_x+w_x} \sum_{y=u_y-w_y}^{u_y+w_y} \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix} \quad (2.96)$$

$$b \doteq \sum_{x=u_x-w_x}^{u_x+w_x} \sum_{y=u_y-w_y}^{u_y+w_y} \begin{bmatrix} \delta I & I_x \\ \delta I & I_y \end{bmatrix} \quad (2.97)$$

yields a final form of

$$\mathbf{d}_{opt} = G^{-1}b \quad (2.98)$$

where \mathbf{d}_{opt} is that displacement vector that causes the derivative $\frac{\partial \epsilon(\mathbf{d})}{\partial \mathbf{d}}$ to be equal to 0. Thus, by calculating the intensity gradients, the optimal displacement vector can be found via the KLT technique. However, because the matrix G must be invertible and because of the limitations inherent in the first order Taylor series approximation, the pixel displacements that are possible to measure must be small.

One means of increasing both the computational efficiency and the spatial distance able to be measured is to subsample an image. One such technique that is able to consistently subsample an image repeatedly and feed information forward and backward from lower resolution to higher resolution images is pyramidal segmentation. Starting with an image of resolution W x H at a level L0, an image pyramid is built by subsampling a smoothed image by a factor of 2 along each coordinate direction at each level. Thus the size of the image at L1 would be W/2 x H/2, at L2 it would be W/4 x H/4, and so on. Figure 2




Figure 2.12: Image pyramid of an ultrasound frame looking at a subject's inferior vena cava created by subsampling segmentation.

demonstrates this process.

Given a point \mathbf{u} in I_1 we want to find its corresponding location in I_2 at $\mathbf{v} = \mathbf{u} + \mathbf{d}$. The corresponding point of $u(u^0)$ on the pyramidal image I^L is \mathbf{u}^L and is defined as

$$\mathbf{u}^L = \frac{\mathbf{u}}{2^L} \quad (2.99)$$

An initial \mathbf{d}^{L_m} is computed at the uppermost pyramid level, L_m , by minimizing its level specific residual function. The final value of \mathbf{d}^{L_m} is then passed as the initial guess for $\mathbf{d}^{L_{m-1}}$ down to the next pyramidal image level, L_{m-1} , to find a value of $\mathbf{d}^{L_{m-1}}$ that minimizes its level specific residual function. This process is continued until the final maximal resolution image at level L_0 is reached, from which a final motion vector, $\mathbf{d} = \mathbf{d}^{L_0}$ is calculated.

Setting $\mathbf{g}^L = [g_x^L g_y^L]^T$ as an initial guess at level L (valid for levels L_1 to L_m), the algorithm seeks to find the pixel displacement vector $\mathbf{d}^L = [d_x^L d_y^L]^T$ that minimizes the

pyramidal residual error function

$$\begin{aligned}\epsilon^L(\mathbf{d}^L) &= \epsilon^L(d_x^L, d_y^L) \\ &= \sum_{x=u_x^L-w_x}^{u_x^L+w_x} \sum_{y=u_y^L-w_y}^{u_y^L+w_y} (I_1^L(x, y) - I_2^L(x + d_x^L + g_x^L, y + d_y^L + g_y^L))^2\end{aligned}\quad (2.100)$$

Starting from the highest pyramidal level, L_m , a guess is initialized as g^{L_m} , and from this guess, an initial motion vector, \mathbf{d}^{L_m} , is calculated. The initial guess and calculated motion are then fed down the pyramid to level L_{m-1} by the function

$$\mathbf{g}^{L_{m-1}} = 2(\mathbf{g}^{L_m} + \mathbf{d}^{L_m})\quad (2.101)$$

Iterating the above function down each level of the pyramidal image yields a final displacement vector

$$\mathbf{d} = \sum_{n=0}^m 2^{L_n} \mathbf{d}^{L_n}\quad (2.102)$$

At every level, the goal is to find the displacement vector that minimizes the matching function. However, rather than search through every possible pixel in an image, a significant computational undertaking for even the smallest and shortest of video clips, one may use the functional relationship between the matching error, $\epsilon(\mathbf{d})$, and the displacement vector, \mathbf{d} to find the optimal displacement vector via the KLT technique at each pyramidal level quickly then passing the information down to the next level (as seen in Figure 2.13).

This algorithm is particularly robust in this use case due the small motions of textured wall patterns caused by the specular reflections from the ultrasound (allowing unique gray patterns to arise) against the black bounds of blood (making demarcation and feature tracking simple). However, to ensure that a robust feature is chosen, the selected point undergoes evaluation through to a simple corner detecting algorithm and slightly alters the point if need be.

Feature selection evaluation

As there is no substitute for clinical experience, no attempt was made to automatically detect and track features.²¹ What was proposed, designed, and created, however, was a check to confirm that a physician-selected point could be tracked across images. This meant that the point had to satisfy some subset of the conditions reviewed in Sections 2.3.1:

²¹At least not successfully.

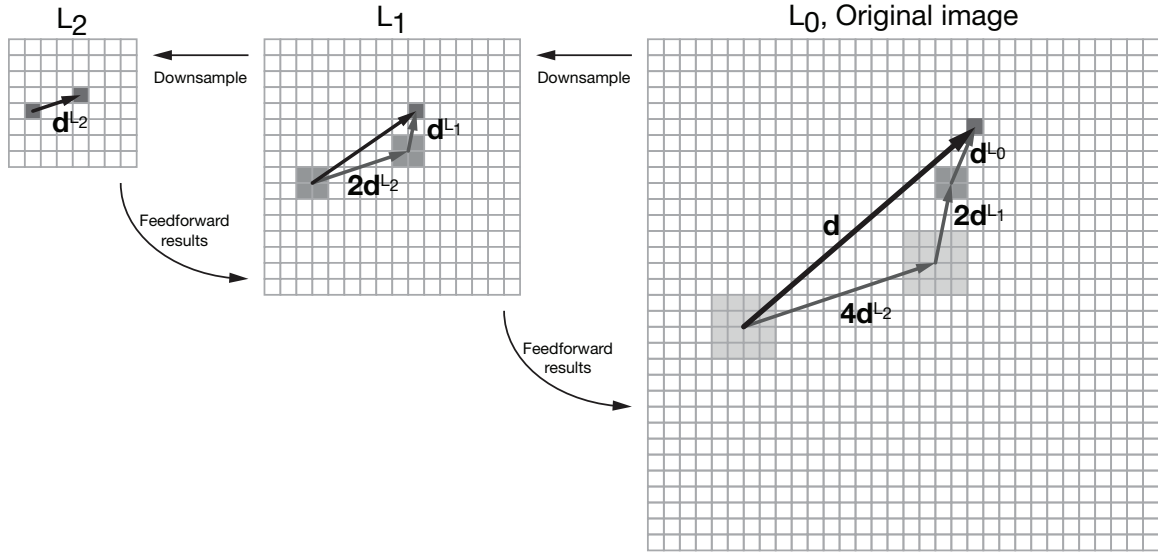


Figure 2.13: An basic example of the iterative process used to track a feature across a pyramidally segmented image. Here an original image, at a level L_0 , is downsampled twice. Once at its highest level, L_2 , initial tracking begins. The results from this tracking are then fed forward to levels L_1 and L_2 , with the results of each, d^{L_m} . adding to final displacement vector, d .

Feature detection and 2.3.2.2: Feature-based methods. Luckily an approach for doing so with satisfying the requirements of the KLT algorithm was suggested by Shi and Tomasi [113] (building off the work of Harris and Stephens [114]). Their technique begins by taking a weighted sum of squared differences (recall equation 2.86), $S(x, y)$, for an image I :

$$S(x, y) = \sum_{(i,j) \in W} w(i, j) (I(x + i, y + j) - I(i, j))^2 \quad (2.103)$$

where $w(i, j)$ is a weighting window, often weighted towards the center. Approximating $I(x + i, y + j)$ through Taylor series expansion and combining equation 2.103 produces

$$S(x, y) = \sum_{(i,j) \in W} w(i, j) (I_x(i, j)x - I_y(i, j)y)^2 \quad (2.104)$$

where again I_x and I_y represent the partial derivatives of I with respect to x and y . This can be written in matrix notation as

$$S(x, y) = \begin{pmatrix} x & y \end{pmatrix} A \begin{pmatrix} x \\ y \end{pmatrix} \quad (2.105)$$

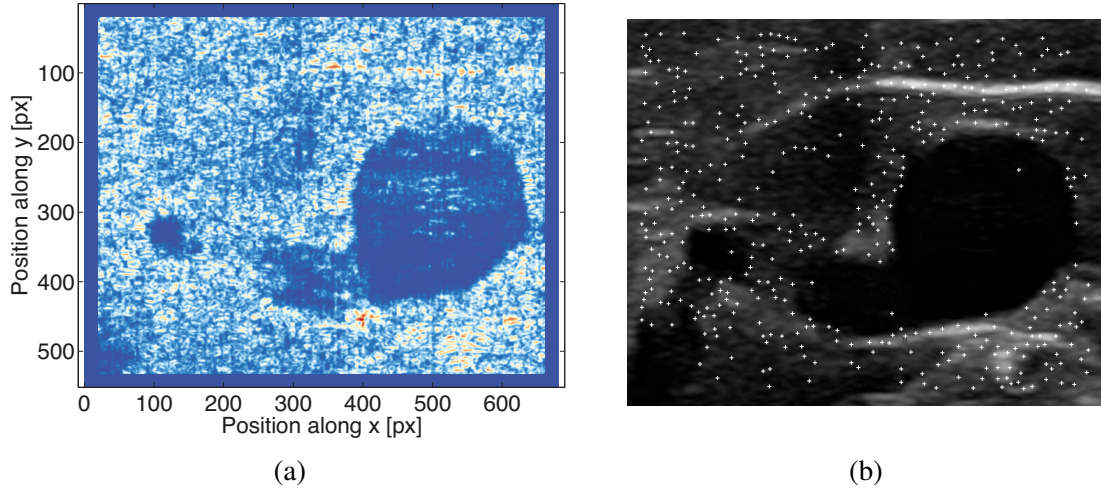


Figure 2.14: Determination of feature selection appropriateness. (a) A map of the eigenvalue pairs values (with red being the highest and blue being the lowest). (b) The top 500 points of maximum value from (a) plotted atop the image from which they were taken. Notice the lack of features seen in the blackness of the blood.

where A is a second-moment matrix

$$A = \sum_{(i,j) \in W} w(i,j) \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix} \quad (2.106)$$

whose resemblance to equation 2.96 should not go unnoticed. A point of interest should produce a large variation in in S in all directions along the vector $\begin{pmatrix} x & y \end{pmatrix}$. Phrased another way, A ought to have two large eigenvalues, λ_1 and λ_2 , at a point of interest. Therefore, by merely examining a weighted second-moment of an image, one can find those points with features distinct enough to track. The limitations placed on G consequently apply to A , making the calculation of minimum and maximum eigenvalues a simple matter.

Figure 2.14(a) demonstrates the implementation of such a technique on a sample image (the same seen in Figures 2.2 and 2.11(a)). From this the maximum eigenvalue pairs can be found and either ranked or thresholded to find features of a sufficient quality. In the case of Figure 2.14(b), the 500 points of maximum value within the image are plotted.

In my implementation, when a user selects a point, the image is subjected to this evaluation and if the chosen point is below some threshold, a spiral search looks for the next available point above the threshold. As nearly every point along a vessel wall is either already sufficiently discernible or within a very short distance of some point that is, this approach was found to be sufficient.

Conclusions

With a good feature tracker and a feature selection verification (and potential augmentation) step in hand, I could now test the algorithm with actual clinical data. My approach was to first evaluate small vessel motions seen in the peripheral vasculature before examining the large changes seen in the inferior vena cava. Though it is ultimately the latter measurement that relates most to the experiments of Chapter 3, the foundations of both are the same and the difference in their validation is merely one of scale and anatomy.²²

2.4 Implementation

Though many ultrasound researchers have the luxury of working with the raw underlying pressure (transduced into a voltage signal at the probe surface), many clinical researchers—most clinicians, physicians, technicians, etc—do not. Available to just about everyone with an ultrasound machine is a means of recording a reconstructed image as it appears on the screen for later use. Because so much important information is lost in the process of transforming the sound waves into pixel intensities, many advanced algorithms employed to elicit more information from ultrasound imaging are unusable for the average clinical researcher trying to go beyond the manufacturer’s in-built functionality. Without the ability to automatically extract information already present in the ultrasound image (such as distensibility of a blood vessel or the strain of a tissue) many innovative biomedical solutions may fall by the wayside.

2.4.1 Example: Arteriovenous fistula mechanical parameters and the development of an open-source ultrasound toolbox

As kidneys fail, so fails the body. Fluids and waste products unable to be removed (as discussed in Section 1.4) accumulate within the body, poisoning it. In cases of severe kidney failure, either acute or chronic, nephrologists may prescribe an extracorporeal waste removal process known as hemodialysis.²³ During hemodialysis (also commonly referred to as simply “dialysis”), blood is removed from the body, pumped through a dialyzer where unwanted solutes (urea, creatinine, etc.) and excess fluid diffuse across a semipermeable membrane via osmotic and hydrostatic pressure, then returned. Long (3-4 hours) and fre-

²²The full software toolbox was designed with fistula optimization in mind, while the feature tracking aspect was meant to track both very small and very large motions.

²³Other renal replacement therapies exist including the kidney transplantation or the closely related techniques of hemofiltration and peritoneal dialysis.

quent (once every 2-4 days) treatments are necessary for those patients with severe chronic kidney disease, especially those advanced into end-stage renal disease. Constant and consistent vascular access is thus, in a very real sense, the lifeline for hemodialysis patients.

Determining an optimal method to achieve long-term, consistent vascular access able to withstand the flows required for dialysis poses unique challenges for clinicians. Currently, there are three main ways to address this: catheter—tubes usually made of silicone rubber inserted into a vein for short-term treatments; grafts—prosthetic tubes usually made of a biologically inert polymer (such polytetrafluorethylene) that connect an artery to a vein; and an arteriovenous fistula—a surgically created connection of an artery to a vein, using the patients autogenous vessels. The latter two types use arterial blood velocities, forces, and pressures to remodel local veins to withstand the large and sustained volumetric flows required for dialysis. As such, the latter two are both created for long-term use, while catheters are employed for emergent or short-term use.

Restricting ourselves exclusively to the upper limbs (whose anatomy can be seen in Figure 2.15) and putting to one side the (still rather significant) role of catheters in dialysis, we are left with two main modes of access: fistulas and grafts. A few common fistulas and grafts their anatomical locations are shown in Figures 2.16 and 2.17, respectively. Generally an initial fistula or graft is placed as distally as possible so that possible failure does not ruin viable proximal vessels still able to transfer blood.

Current guidelines recommend autogenous arteriovenous fistula as the preferred vascular access for patients undergoing hemodialysis based on studies showing longer patency and few complications [115, 116, 117]. These fistulas, however, often fail to mature at alarming rates (anywhere from 10-70%) [116, 117, 118, 119, 120]. Such failures interrupt life sustaining dialysis treatments and lead to procedures and sometimes hospitalizations for revision or new fistula creation, adding to the overall cost of care, increasing the medical risks, and causing unnecessary stress on patients. To address this issue many clinicians have turned to non-invasively measuring certain morphologic parameters including artery and vein diameter, their respective distensibilities and volumetric flows [121, 122, 123, 124] via ultrasound as predictors of fistula maturation, aimed at finding anatomical, physiological, and mechanical feedback systems regulating the patency of fistulas.

Many of these types of measurements must be explicitly made at the bedside, with a particular machine, or processed manually via proprietary software, complicating the study of these metrics by making them less available to interested clinicians.

The goal is to measure predictive factors for each patients fistulas likelihood of success or failure by developing a rigorous and openly available ultrasound software tool for fistula mechanics quantifying the interactions of vessel elasticity and blood-wall shear rate

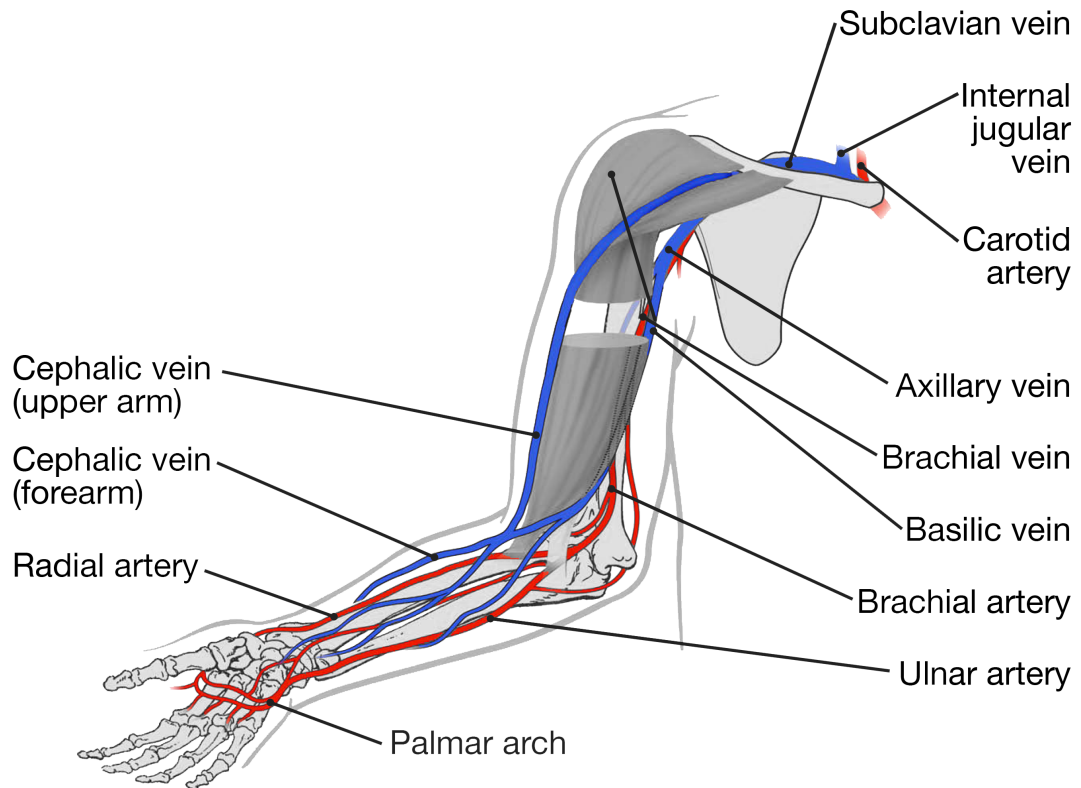


Figure 2.15: Anatomy of the upper limb. Several blood vessels relevant to arteriovenous fistulas (brachial, radial, and ulnar arteries; axillary, brachial, basilic, and cephalic veins) are highlighted. (Modified from the *Atlas of Dialysis Vascular Access* for educational purposes.)

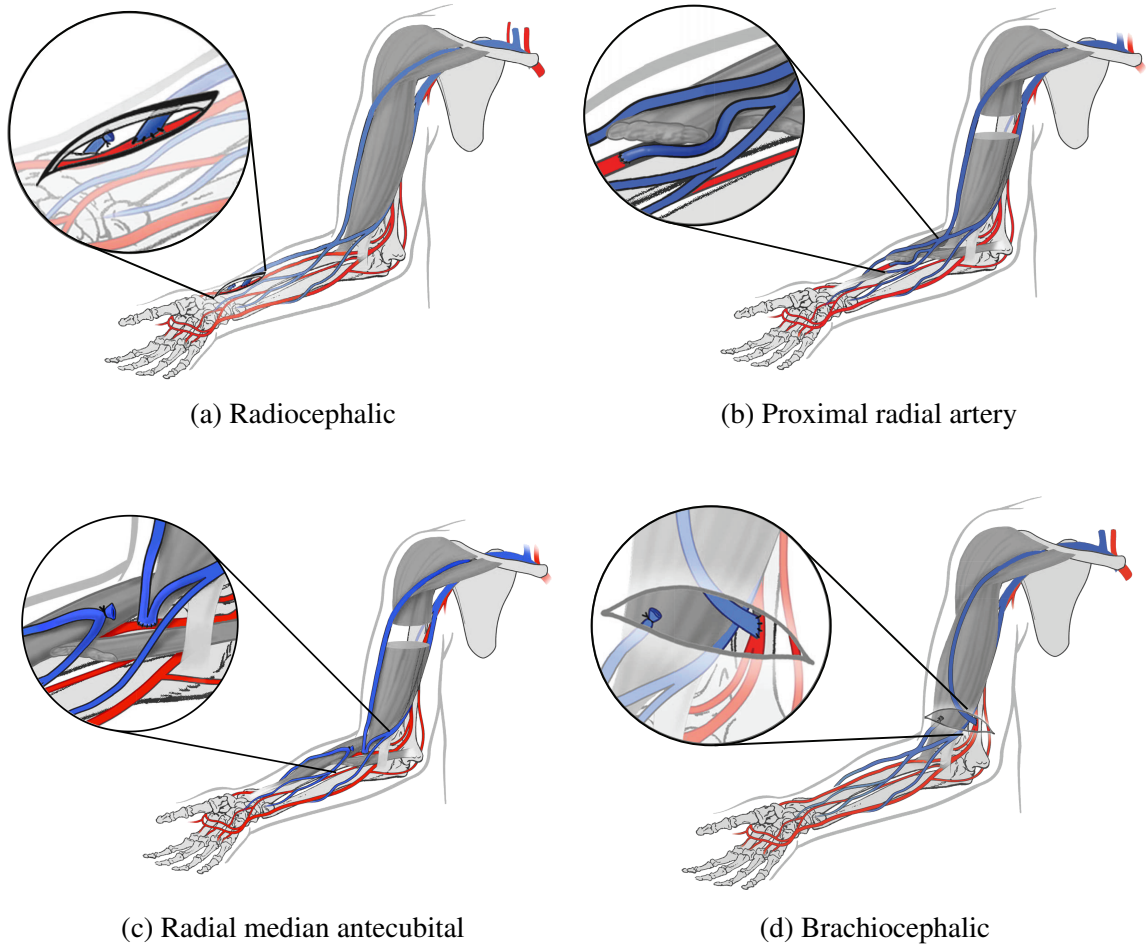


Figure 2.16: A few major types of arteriovenous fistulas. These fistulas are formed by the pairing for the the (a) cephalic vein and radial artery, (b) perforating branch and proximal radial artery, (c) median antecubital vein and radial artery, and (d) cephalic vein and brachial artery. (Modified from the *Atlas of Dialysis Vascular Access* for educational purposes.)




Figure 2.17: Two types of peripheral graft-based vascular accesses: (a) forearm loop and (b) upper arm. In both cases a synthetic tube redirects arterial blood to the venous side to modify the size, shape, and physical properties of the veins. (Modified from the *Atlas of Dialysis Vascular Access* for educational purposes.)

within the areas of interest (the artery, the vein, and post-arterial anastomotic segment after surgery). This software toolbox was designed to evaluate stored ultrasound images in the Digital Imaging and Communications in Medicine (DICOM) standard for parameters including distensibility, strain, and blood-wall shear rate by using verified imaging processing techniques.²⁴ The software functionality, measurement results, and its application for optimization of vascular access surgery are presented.

2.4.1.1 Overview of a software toolbox

The software toolbox is designed around the graphical user interface (GUI) implemented in MATLAB that allows a user to open files, select regions of interest, analyze ultrasound images for parameters, and save data (see Figure 2.18). MATLAB was chosen as the environment and programming language because of its broad access within research environments, vast in-built functionality, and its large community of user. This toolbox is available for implementation and modification by researchers.

Upon initializing the software, a user is prompted to browse for and open an ultrasound image or video that is either saved in a DICOM format or as a MATLAB array. Once the image or video is loaded (if it loads a video, the first frame will be displayed), the

²⁴An additional function to reconstruct the three-dimensional geometry of vessels is also being worked on at the time of publication and its preliminary results are included in Section 2.4.1.5: Image segmentation and geometry reconstruction.



Figure 2.18: Main GUI of the software toolbox developed to aid in fistula maturation research. Its main features from this window include (a) viewing the color histogram of the image from which (b) a user can manually set brightness and contrast parameters, (c) a zoom-to-fit function to get the most use out of the space, (d) overlaying a grid on top of the image, (e) the ability to change the colormap, (f) a save function for the marked-up ultrasound image in the window as either (g) a single frame or the entire video, (h) the ability to edit or export data made by the user during a session, (i) importing a new image to work with, (j) displaying the current frame number and the total number of frames, (k) applying either a median or Kalman filter across the frames to minimize certain types of noise in the images, (l) setting up the image properly (such as inputting or calibrating the distance between pixels, determining the number of pixels to be analyzed during strain measurements, establishing the bounds for accumulated strain, inputting blood pressure values, specifying the frame rate, or asking for help), (m) analyzing the pulsatility, strain (either total image or vessel wall specific), or blood-wall shear rate of the image, (n) manipulating the image for analysis (such as cropping (either manually or automatically), mapping pixel intensities for a given perimeter, measuring point-to-point distances, or removing frames), (o) specifying regions of interest, (p) displaying mean and standard deviations of certain measurements within regions of interest, (q) displaying the pixel intensity values at the cursor, and (r) scrolling through each of the video frame by frame.

user has a number of manual controls they can take advantage of to make the viewing of an image more suited to a particular persons needs including brightness, contrast, and color modifications. If the user selects a video, they can scroll through frame by frame. Users are also able to make many simple measurements at any time such as determining individual pixel values, measuring distances between points of interest, cropping to a region of interest, annotating the image with basic shapes, and saving frames of interest for future examination or display.

Where this software toolbox excels for vascular access researchers is in providing measures of three key parameters: (1) distensibility tracking vessel diameter changes over time; (2) strain measuring the deformation with reference to a desired geometry; and (3) wall shear rate the rate at which fluid shears across the vessel wall. Each of these measurements is achieved through semi-automated digital image correlation methods that relate changes in pixel intensity arrays occurring within specified regions of one frame to another. As these three mechanical parameters are thought to be critical for predicting AV fistula maturation and as their measurement through these image-processing techniques has been previously demonstrated [125, 126, 127, 128, 129, 130], the primary purpose of this software was to thus give investigators the tools by which to measure them.

Furthermore, an emphasis was placed on local feature tracking throughout whole images and within cardiac cycles to give researchers sufficient granularity to target individual aspects of the maturation process. For example, with this toolbox a researcher could correlate the systolic wall shear rate on the inner and outer walls of the post arterial anastomotic segment and correlate its trend with the increasing size of the vein throughout maturation. Specific use cases of this toolbox are myriad and we will endeavor here to show the robustness of the feature set.

2.4.1.2 Distensibility

To measure the distensibility of a vessel, the user is first prompted to select two points of interest. These points are then tracked using the Kanade-Lucas-Tomasi feature-tracking algorithm [111, 131] previously discussed, native to the MATLAB image processing toolbox and OpenCV. This feature tracker looks for spatial intensity gradients within images to guide the search for the maximum frame to frame correlation for a subset of pixels, and thus quickly and accurately locating where a point of interest has moved. Because the unique patterns caused by the specular reflection at the vessel wall provide sufficiently unique and discernible detail, this algorithm should be useful and robust at tracking the diameter changes. Having tracked the distance between these two points over the duration of the cine loop, the user is then presented with a graphical representation of distensibility over time.

The vessel diameter can be displayed as a measure of absolute distance, relative distance, normalized as a percent of the maximum value (to yield what is more commonly referred to as “collapsibility”), normalized as a percent of the minimum value (“distensibility”), or normalized to the range (to see the relative shape of each waveform).

Figures 2.19 demonstrate the results of the feature tracking algorithm to measure the diameter of a single vessel measured with ultrasound over time. (In this case, it is the same vessel as that through which the A-line passes in Figure 2.2).²⁵ We can see from simple graph a number of interesting traits thus far not reported in the literature to any great degree, though their presence is expected from physiology. One is the visibility of a respiratory contribution—increasing and decreasing the diameter of the vessel—that at first may appear as a type of drifting error. However, the low frequency signal corresponds well to a subject’s respiration, disappearing in cases of inspiratory and expiratory holds.²⁶ Another feature of note is an appearance consistent with arterial flows and pressures: a sharp rise accompanied by a decay with signs of a dicrotic notch. One final characteristic to emphasize is the clarity with which the heart rate and respiratory rate can be measured from this, something impractical (if not impossible) to attempt to do manually.


Multiple paired point selection functionality is also available for those researchers who wish to examine the relationship between different geometric parameters simultaneously, such as an artery and vein simultaneously or the circumferential strain measured by several points along a vessel wall.

2.4.1.3 Strain


Extending the basic pixel- and region-tracking function from a single pair of user selected points to an automatically generated grid of pixels across the image and incorporating optical flow techniques outlined in Section 2.3.3 Image transformation, the software is able to determine the overall changes in distance over time throughout the region of interest. In so doing, a user can measure the strain present in the region in one of two ways: frame-to-frame where all change in one frame is scaled relative to the frame that came before it or cumulatively where a single reference frame is chosen to scale all subsequent frames against. In both analyses, the user selects what fraction of the pixels they wish to include to measure (ranging for 1-100% in increments of 1%) and a uniform grid of pixels is selected from the reference frame to be tracked in the subsequent frame(s). Lateral, axial, or the overall magnitude of strain can be selected by the user to suit their needs. Strain maps themselves can be saved as single frames, video clips, or arrays of data for future process-

²⁵In fact, if one were to squint just enough at Figure 2.4, one would see the same diameter change.

²⁶Recall the effects of intrathoracic pressures on the cardiovascular system discussed in the first chapter.



(a)



(b)

Figure 2.19: The (a) diameter and (b) distensibility of the vessel measured in Figure 2.2 over time. On the left-hand y-axes the absolute measures of diameter and distensibility can be seen, while on the right hand y-axes the relative measures of diameter (from minimum) and distensibility (across range) are shown.




Figure 2.20: Strain map development of a radiocephalic fistula over time. The magnitude of strain is presented here (a) just after the initiation of a cardiac cycle, (b) halfway to peak strain, and (c) peak strain. The pixel intensity represents the magnitude of strain and local maximums can be found at both the arterial wall and one edge of the post-arterial anastomotic segment. (These images were produced by subsampling the original image by 20%.)

ing. Figure 2.20 shows the cumulative strain magnitude map of the radiocephalic fistula shown in Figure 2.18 taken over a single cardiac cycle. Both the artery and the post-arterial anastomotic segment experienced large local strains during peak pulse pressure.

2.4.1.4 Wall shear rate

The wall shear rate measurement the software toolbox provides is based on speckle decorrelation method that has shown efficacy in determining flow velocities near vessel walls [127, 127]. The method works by assuming that as scatterers pass through the ultrasound imaging beam, the degree to which they correlate from one image to another is a function of the distance that the scatterer moves. Thus correlation (and its inverse, decorrelation) becomes a measure of distance. Coupling the measure with the sampling frame rate that is, by knowing the temporal distance separating the spatial correlation one is able to infer a velocity. One drawback to this technique is its reliance on high sampling frequencies to measure high velocities. Because typical cine loop frames rates for commercial ultrasound machines are usually between 10-50 Hz, it would be extremely difficult to consistently measure full flow profiles even within peripheral vasculature. However, near the vessel wall, where the flow velocities are often their lowest, this technique can excel over alternatives such as Doppler flow measurements. Furthermore, since the forces that affect mechanotransduction and ultimately shape the vessel occur at the wall, it is of the utmost importance to reliably measure the flow velocity and its gradient, the wall shear rate.

Should a user opt to use the wall shear rate, they will first be prompted to select a point




Figure 2.21: Blood velocity profile in a brachial artery at peak systole from both a decorrelation measurement and as an ideal Poiseuille flow.

along the wall and then another point perpendicular to the wall. By giving the user control in this selection process specific regions and flow orientations can be targeted. Once the decorrelation algorithm has been run, the user will then be prompted to input the functional parameters of their own correlation vs. distance curve, specific to their machine of use, to get a quantified flow velocity and wall shear rate (as described by the methods of Park et al. [128]). If a user does not have such a curve, they will be presented with a figure of correlation vs. distance and warned that without an input function this graph should be treated as a qualitative assessment of the bloods velocity near the wall.

Figure 2.21 demonstrates an example of a peak brachial artery flow velocity profile at the vessel edge measured via speckle decorrelation and an ideal Poiseuille velocity curve (calculated from patient dimensions and Doppler volumetric flow measurement). The wall shear rate is calculated by fitting a line across the first three velocity points from the edge. In this case a wall shear rate within approximately 10% of an ideal model was found and is within the standard range of peak wall shear rate values for the brachial artery [132, 133], indicating a reliability sufficient for pilot clinical evaluation.

2.4.1.5 Image segmentation and geometry reconstruction

At the time of this publication, I have a function that can plug into the toolbox and attempt to reconstruct a vessel's three-dimensional geometry. Because of the difficulty of the procedure to acquire images that produce this geometry and the somewhat tenuous assumptions necessary in its reconstruction, it has not yet been included into the distributed toolbox for

popular adoption. I encourage those researchers interested in further information, code, or demonstrations to contact me. For this work, I merely present a brief overview of the technique.

The procedure begins by obtaining a short-axis view of the vessel of interest via an ultrasound transducer. The transducer must then be moved along the longitudinal length of the vessel while recording sequence of B-mode images. The steadier the longitudinal motion, the better.²⁷ Once the sequence of images is obtained, the edge of the vessel must be found. While there are many excellent algorithms for doing so (as mentioned in Section 2.3.1), ultrasound images are particularly prone to produce an overabundance of “edges.” I employed a simpler method of simply segmenting the images via grayscale thresholding. In this way the black interior of the vessel can be easily separated from the grays and whites of the surrounding tissues. The pixels along the boundary may then be treated as the edge for each frame. An enclosed surface may then be fit across a vessel via a fitting function of one’s choice. (My particular implementation sought to minimize the sum-of-squared differences between the splined interpolant and the surrounding points.) In the case of a fistula, where a branching occurs, two separate vessels must be created and merged. The point cloud, splined, and final geometries of a radiocephalic fistula (the same as that shown in Figure 2.18) can be seen in Figure 2.22.

Several limitations exist to such an approach. The most obvious and detrimental is the assumption that the ultrasound beam is kept constantly perpendicular to the cross-section of interest. In practice this is difficult to the point of impossible to achieve. To compensate for this, one could use the probe’s orientation in space with respect to a patient’s anatomy to correct for projective geometry distortions (not heretofore discussed). Several excellent techniques exist to implement such a feature [134, 135, 136, 137, 138], but for time and scope constraints were not implemented in this work. Extremely pulsatile vessels (such as the inferior vena cava) could cause temporal changes to become spatial distortions and thus produce wavy vessels.²⁸ Small vessels, too, pose difficulties as the pixel intensity values between the interior and exterior of the vessel are not, if I may be excused just this one bit of wordplay, so clearly black and white.

²⁷Trained hands or automated methods are prerequisite.

²⁸However, an enterprising researcher could use this pulsatility to their advantage. If one splined minimum peaks to minimum peaks to minimum peaks and maximum peaks to maximum peaks, one might be able to reconstruct both the minimum and maximum vessel geometries. Furthermore, the relaxation of the vessel from the maximum to the minimum could also be used to back out elastic properties, stresses, and strains. Alas, such work is for another dissertation by another researcher.

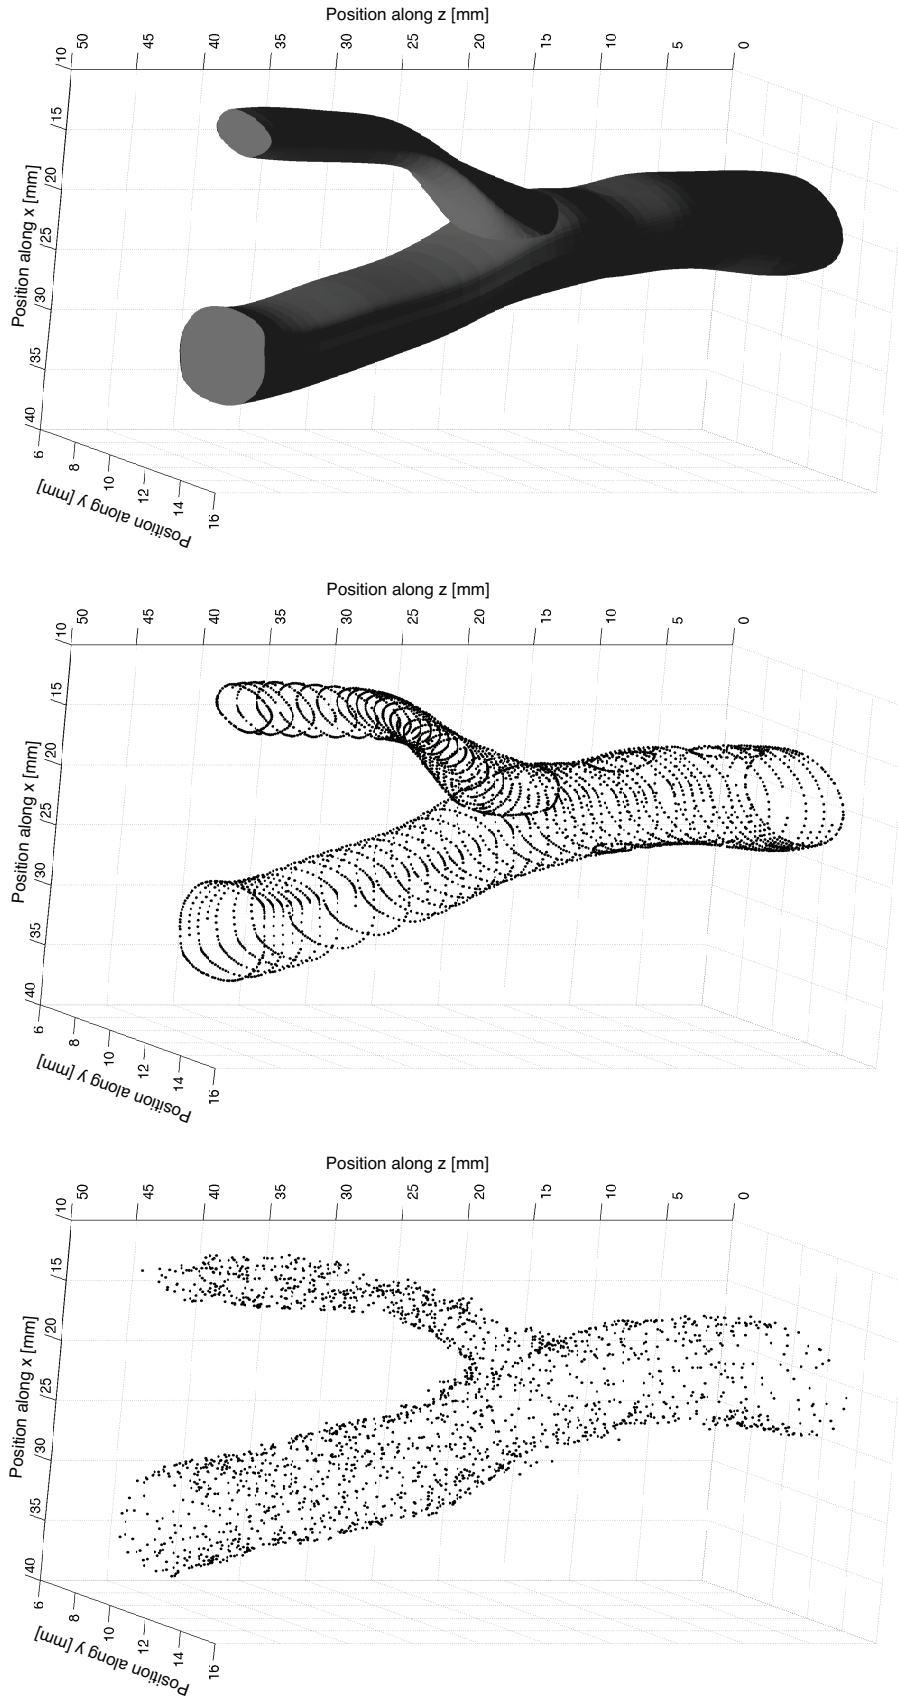


Figure 2.22: The initial segmented vessel geometry (left), a smoothed and splined edge geometry (middle), and the final solid three-dimensional geometry (right). Only 1% of available edge pixels are displayed in the initial segmented image to make its properties clear. Similarly, every other image slice is shown for the smoothed and splined geometry. The final three-dimensional geometry here was used for some of the fluid simulations reported in Figures 2.23-2.25.

2.4.1.6 Theoretical, simulated, and measured flows

Flow through the fistula is one of the key determinants of patency. Pressures, shear stresses, shear rates, and velocities all serve interweaving and dynamically varying purposes in the vasculature in this region. The anastomosis should receive a significant influx of arterial flow to remodel the vasculature in response to the higher flows [139]. Elevated flows are known to increase nitric oxide mediated vasodilation and the long term exposure of elevated shear stresses to the endothelial cells of the wall spur adaptive remodeling of the vessel wall at a cellular level.²⁹ An imbalance of these facets (too much of one, too little of another, etc.) can cause fistulas to fail to mature.

Guidelines exist to aid vascular surgeons in the creation of fistulas and I have developed an image processing toolbox to help with the evaluation of fistulas, but the two approaches must be bridged by some basic intuition of the flow patterns caused by these combined vessels. To that end, I conducted a series of simulations to model flow within these vessels to determine what factors and which regions might be of most clinical interest when using my fistula toolbox.

The simulation performed was of the simplified patient-specific geometry reconstructed via the steps of Section 2.4.1.5 and presented in Figure 2.22. To simplify the situation immensely, these simulations assume a constant flow from the artery (of 1200 mL/min) and a proportional mass transport boundary condition at arterial and venous outflow (of 40% and 60%, respectively), found by via Doppler flow measurements from the actual patient. Unlike water or saline, blood is a non-Newtonian fluid whose viscosity depends on the velocity gradient. I chose to work with a modified Power Law first proposed by Jozwik et al. [142] for the dynamic viscosity term, μ , as it has been done previously for similar arteriovenous fistula modeling [143]. Its definition is as follows:

$$\begin{cases} \mu = 0.554712 & \frac{\partial V}{\partial y} < 10^{-9} \\ \mu = \mu_0 \left(\frac{\partial V}{\partial y} \right)^{n-1} & 10^{-9} \leq \frac{\partial V}{\partial y} < 327 \\ \mu = 0.00345 & \frac{\partial V}{\partial y} \geq 327 \end{cases} \quad (2.107)$$

Three results relevant to the measuring capabilities of the software toolbox were looked into and are reported here. The first is the velocity profile of Figure 2.23 showing regions of high velocity in the post-arterial anastomotic segment and lower velocities in the outflowing artery. This agreed with the results found from the Doppler ultrasound of this patient and

²⁹Specifically, the apical cell surface deforms in the direction of flow causing cytoskeletal reconstruction and activating several key signaling cascades in response to acute and chronic release of nitric oxide. [140, 141]




Figure 2.23: Velocity streamlines across the simulated fistula.

the values obtained are in line with those found elsewhere in the literature. The rotational flows found were also in line with expectations given what we discussed in Section 1.3. Figures 2.24 and 2.25 show the distributions of wall shear stress and the wall shear rates across the fistula respectively. As is expect from intuition and the reports of others, large shear stress concentrations occur in those regions with the highest bulk velocities. The peak wall shear rates occur along areas of the outer and inner edge of the post-arterial anastomotic segment, in line with previous reports.

2.4.1.7 Clinical evaluation

A pilot study was conducted twenty-four patients who received an arteriovenous fistula was conducted to demonstrate the feasibility of clinical implementation of this software to evaluate these mechanical parameters. Subjects were consented before surgical creation of their fistula. As part of their standard of care, an ultrasound evaluation of the fistula was performed approximately one week and five weeks post operatively, and it was during this time that long- and short-axis views (longitudinal and transverse views, respectively) of the artery, post-arterial anastomotic segment, and vein were obtained by a trained technician using a commercially available ultrasound machine with a vascular access probe (need specifics). The inflow and outflow arteries and veins were each imaged approximately 10 cm from the surgical anastomosis.




Figure 2.24: Wall shear stress across the simulated fistula.




Figure 2.25: Wall shear rate across the simulated fistula.

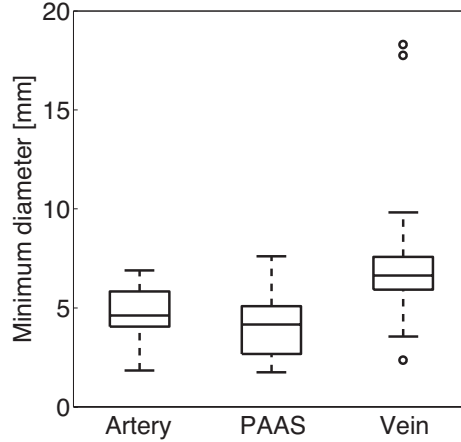


Figure 2.26: Minimum diameter of the inflow artery, post arterial anastomotic segment (PAAS), and the outflow vein of measured during an initial evaluation ($n = 24$ patients). The minimum diameter is reported here as the parameter of interest, distensibility—the amount a vessel expands from a baseline—relies on it. While the distributions of the vessel sizes are fairly normal, a few outliers exist in the veins with a pair of measured veins very large and a single small vein. The patients to whom these outliers belong were removed from the follow-up examination evaluation (seen in Figure 2.28).

Figure 2.26 shows the minimum diameter of the inflow artery, post-arterial anastomotic segment, and the outflow vein during the initial evaluation (approximately one week after surgical creation). As would be hoped for, the vein was found to be on average the largest vessel. Three notable outliers are present in the vein examination, with two subjects having an overly large vein and one subject having an abnormally small vein. Though the large

Figure 2.28(a) and (b) show the changes in vessel size and distensibility, respectively, for the artery, post-arterial anastomotic segment, and vein. These preliminary results indicate that over maturation, the artery tends to become more pulsatile (going from 2.67% to 4.72%)—agreeing with trends observed by others through different means [144]—and the vein tends to get larger (from an average value of 6.18 to 8.10 mm), as is expected. The post-arterial anastomotic segment experiences the widest variation among patients upon first inspection, but after five weeks, the variability in the distensibility dropped significantly, with these patients’s values settling to an average of 4.2%. These trends are in line with our current understanding of the vascular remodeling that fistulas undergo, indicating that the vein becomes larger and more like an artery while the artery tends to become more compliant, like a vein. In between, their junction, the post-arterial anastomotic segment takes on attributes of both.

An estimate of the elastic modulus of the vessels could also be found by first positing that the distensibility is equivalent to a mean circumferential strain, ϵ_θ , via the Cauchy

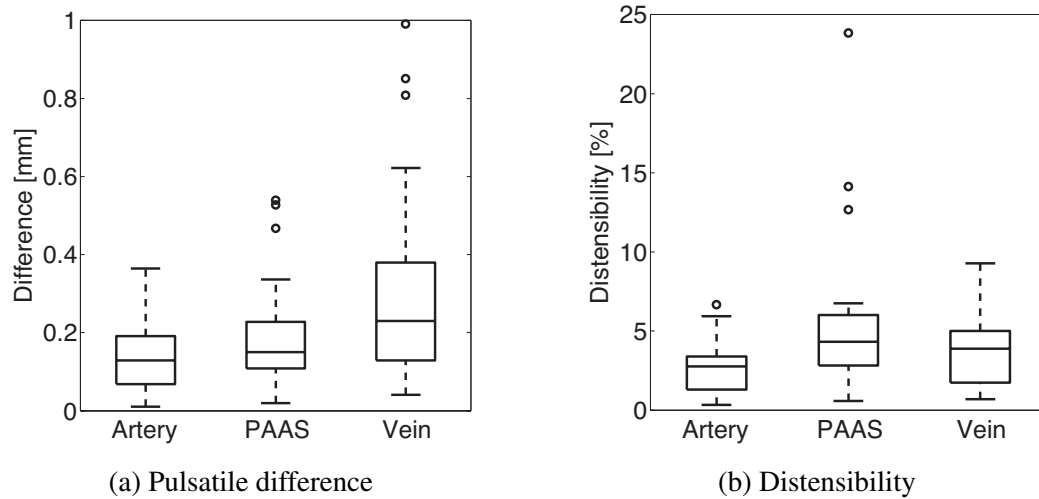


Figure 2.27: Distribution of (a) pulsatile differences and (b) distensibilities for the inflow artery, post arterial anastomotic segment, and outflow vein measured during an initial evaluation (n = 24 patients).

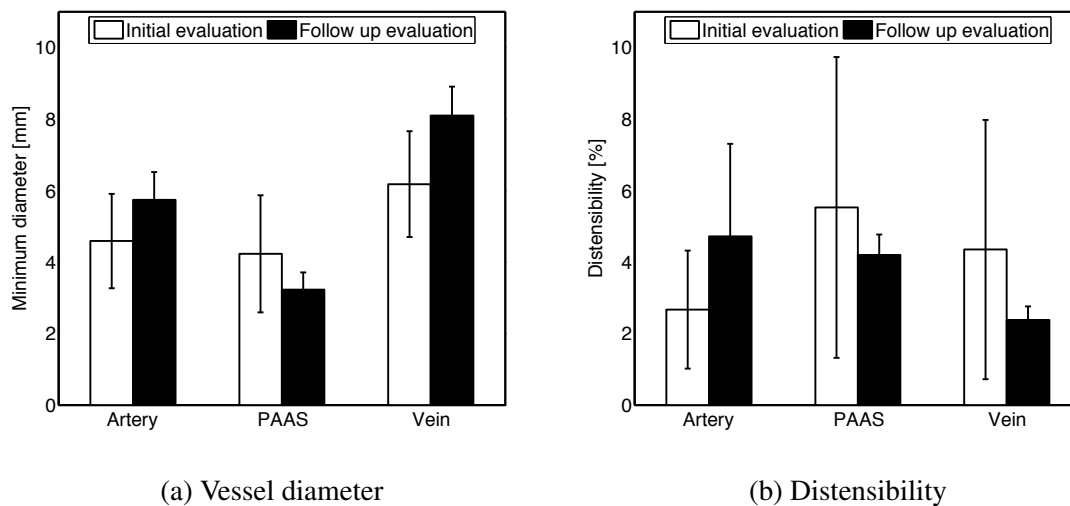


Figure 2.28: Clinical data relating the change in (a) vessel diameter and (b) distensibility measured during an initial evaluation (approximately one week post-operation) and a follow up examination (typically one to two months post-operation). Note the increased size of the vein and the increased pulsatility of the artery in the follow up evaluations as compared to the initial evaluations.

strain definition, such that the strain at any given time, t , is

$$\epsilon_{\theta}(t) = \frac{D(t) - D_{min}}{D_{min}} \quad (2.108)$$

where $D(t)$ is the diameter of the vessel at a point in time and D_{min} is the minimum diameter of the vessel. Assuming the walls of the vessels are much smaller than the diameter itself, the mean circumferential stress, $\sigma_{theta}(t)$, is then

$$\sigma_{\theta}(t) = \frac{P(t)D(t)}{2h} \quad (2.109)$$

where $P(t)$ is the pressure at a point in time and h is the wall thickness. Such an approach has been used elsewhere effectively [145]. It is important to state here that the assumption of thin walls used is not necessarily valid. Unfortunately for this work, actual wall thickness was not measured and thus must itself be estimated (as is shown below). Inaccurate uniformity and variability must therefore be ever present in our minds for the rest of this conversation. However, no stenoses or oddly shaped cross-sections were observed. Hence I feel comfortable inferring that the vessel walls of these patients were within a normal range, thereby satisfying the thin-walled condition, and presenting this examination. I estimated the wall thickness of the vessels by synthesizing the results of [146, 147, 148, 149, 150] to reckon that the ratio of the intima-media to the lumen diameter was equivalent to approximately 0.05 (on the higher end of the estimates provided by [146] by extrapolating the relationships found for patients with peripheral artery disease, hypertension, and high blood pressure—all sharing some commonality with the end-stage renal disease patients of this work). Much of the nuance of the possible results is unfortunately lost here for the investigation to happen at all.

The pressures during the cardiac cycle are present and ever changing producing a time varying stress-strain relationship. Two well defined points exist, though, to aid in this analysis: peak systole and end diastole. Systolic and diastolic pressures correspond to the maximum and minimum distention of the vessels, respectively, and allow a simple linear estimate of the elastic modulus to be calculated. Because of the nonlinearity mechanics of vessel walls [151], the estimates that follow must be viewed more as a qualitative assessment of arteriovenous fistula maturation more for the benefit of clinicians than as a quantitative determination of vascular mechanics.

The elastic moduli of ten patients' vessels is presented in Figure 2.29. No statistically significant results were observed between the artery and post arterial anastomotic segment before and after ($p = 0.387$ and $p = 0.454$ using a paired, two-tailed t-test). The outflow-




Figure 2.29: Compliance of the artery, post-arterial anastomotic segment, and vein during the initial evaluation and the follow up examination. (Data presented is of ten patients with complete records.)

ing vein’s elastic modulus, however, was statistically distinct from the initial evaluation to the follow-up evaluation ($p = 0.0112$), suggesting that the vein become less compliant as the fistula matures. Whether this stiffening is an artifact of increased local venous blood pressure or vascular remodeling remains to be seen.³⁰

2.4.2 Example: Continuous inferior vena cava tracking

As previously mentioned, one of the most promising techniques to rapidly gauge fluid responsiveness has been ultrasound measurement of respiratory induced changes in the inferior vena cava (IVC) diameter. Having been validated across a variety of conditions (sepsis [152], trauma [56], heart failure [53], dialysis [153]) and a variety of patient types (spontaneously breathing [72] and mechanically ventilated patients), the collapsibility of the IVC has been shown to reliably guide fluid therapy [154, 69, 70, 71, 155].

The technique works because as one breathes the shifting diaphragm causes intrathoracic pressure gradients to rise. In spontaneously breathing people, the diaphragm pulls down during inspiration, creating a transient but significant negative intrathoracic pressure change. This decrease in intrathoracic pressure results in an increase in venous return to the heart via the vena cava. This relationship is reversed in the case of positive pressure ventilation, where a positive intrathoracic pressure accompanies inspiration. As a very

³⁰Though average systolic and diastolic pressures during the initial evaluation were on average higher than those during the follow-up examination, there was not a statistically significant difference observed. Presented as the mean(standard deviation) of systole/diastole the initial measurements were 141(28)/70(11) mmHg while the follow up measurements were 131(39)/64(10) mmHg.

compliant blood vessel, the IVC is affected by these changes in flow and will collapse to various degrees under negative pressures (inspiration during spontaneous breathing, expiration under mechanical ventilation) and expand under positive pressures (expiration during spontaneous breathing, inspiration under mechanical ventilation). The extent of collapse (or distension) is a function of volume status and of the amount of blood returning to the heart (venous return) [3, 156] and thus can help guide fluid resuscitation.

Respirophasic IVC changes are a dynamic measure of right ventricular preload. In spontaneously breathing patients a collapse of greater than 40% indicates fluid responsiveness [72] the patient will increase cardiac output if provided fluid [157]. In mechanically ventilated patients an increase IVC diameter of 12-18% similarly indicates responsiveness [154, 70]. Such a measurement can be taken quickly and easily at the bedside via ultrasound (Figure 1.14), providing actionable data in the time it takes to take a breath. Because of the relative ease in taking the measurement, its predictive power of fluid responsiveness, and the increased utilization of point-of-care ultrasound systems, the caval index as a guide for fluid therapy has become increasingly popular among clinicians.

Limitations to a physician-based measurement technique exist, however, including relying on rough estimates of collapse by physicians without measurement and inconsistent collapse measurements stemming from measuring different points along the vessel. Moreover, currently no tool is clinically available that can measure the diameter of the IVC continuously during the performance of the ultrasound.

To provide such a tool, we developed a simple point-tracking algorithm based on the paired combination of the Kanade-Lucas-Tomasi feature tracker [111, 131] and pyramidal segmentation. Working with ultrasound videos in the DICOM (Digital Imaging and Communications in Medicine) standard, this simple algorithm allows a clinician to manually select points on either side of the vessel to track over the course of the video. The distance between these two points is then calculated and the respiratory variations in the diameter of the IVC can be calculated for the caval index. With this technique, we have created a tool to investigate the continuous waveform of the IVC diameter, opening many avenues for future research, including the recently reported cardiac collapsibility index [158].

In this section I outline a clinical experiment designed to test the accuracy of the feature tracking algorithm discussed previously to the task of tracking the diameter of the inferior vena cava via ultrasound. The results of the algorithm are compared against specially trained physicians and the automated measure was found to be accurate to within the level of variation between clinicians themselves.

2.4.2.1 Experimental design

Fifty spontaneously breathing subjects had their IVCs examined while in a supine position. A physician with training in ultrasound imaged the IVC subcostally in a subxiphoid longitudinal view using a phased array or curvilinear ultrasound probe. Ultrasound was performed using a commercially available device (Mindray M7, Mindray North American, Mahwah, NJ). A total of 64 B-mode cine loops 10-15 seconds in length were saved in the DICOM format for post-clinical evaluation. Though many sonologists have used M-mode imaging to calculate IVC collapsibility and distensibility, it has been demonstrated that respiration results in caudal displacement of the IVC, often leading to inaccurate M-mode based measurements [159, 160].

Two trained ultrasound physicians then measured (via Escape Medical Viewer, Thessaloniki, Greece) the caval index by identifying a point of maximal collapse (minimal diameter, peak of inspiration), approximately 2 cm caudal to the hepatic vein inlet (as suggested by Wallace et al. [159]), then using frame by frame manual analysis to find the corresponding point of maximal diameter (peak of expiration). The caval index, CI , for spontaneous breathing is here defined as the difference between the maximum, D_{max} , and minimum, D_{min} , diameter normalized to the maximum diameter, represented as a percent

$$CI = \frac{D_{max} - D_{min}}{D_{max}} \quad (2.110)$$

A wide variety of inferior vena images (see Figure 2.30) were measured to test the limitations of the automated measurement.

Using the same landmarks, points were selected by a researcher using the software that were then tracked frame-to-frame via the algorithm described in the previous section. The minimum and maximum points within a clip were detected via peak detection and the caval index for each clip was calculated. The results obtained from the clinicians and the software were compared.

2.4.2.2 Statistical analysis

Descriptive statistics including means, standard deviations and diagrams were found to visually inspect the data. Intraclass correlation coefficients (ICC) (2,1) [161] were calculated to determine the level of agreement and consistency between the physician and automated measurements. The method described by Bland and Altman was used to analyze the bias and precision between automated and physician measurements and to determine the limits of agreement by defining the mean and standard deviation (SD) of the differences. The limits of agreement were defined as the mean difference $\pm 2(\text{SD})$ of the differences. Data




Figure 2.30: Varieties of inferior vena cava images. A point tracking system must be robust enough to handle the substantial variety of ultrasound images that arise from patients.

are presented as mean(SD) and where appropriate, 95% confidence intervals are reported. For all calculations, p levels of ≤ 0.05 were considered statistically significant. MATLAB was used for all data analysis.

2.4.2.3 Results

Of the 64 initial images, three were deemed not possible to measure by the clinicians (lacking the necessary landmarks or suffering from significant cylinder tangent throughout a clip). An additional four were not able to be measured by the algorithm due either to consistent software crashes across many versions of MATLAB (two cases) or images of too poor in quality to track (two cases, discussed later). Therefore a total of 57 images from 47 patients were included for analysis.

One result observed through the automated measurement that is difficult to discern manually by a clinician is diameter changes corresponding to both respiratory and cardiac cycles. Figure 2.31 demonstrates the effects of both cycles in a spontaneously breathing subject. The raw signal in the top graph of Figure 2.31 is the result straight from the KLT algorithm described. To isolate the respiratory signal, this raw signal was low-pass filtered with a corner frequency of 0.6 Hz. Subtracting the respiratory signal from the raw signal resulted in what I believe to be signal from the cardiac cycle (Figure 2.31 (bottom)). Only recently have researchers begun to explore the potentially useful metric of the cardiac collapsibility index (CCI) in the IVC [158]. To my knowledge no researchers have previously presented evidence of the existence of CCI variability (CCIV) as shown in the bottom half of Figure 2.31.

For the results that follow, what is presented are the maximum and minimum diameters of the raw signal. As the physicians had no way of knowing the individual contributions of the cardiac and respiratory cycles, the accuracy of their individual measurements could not be made.

Figures 2.32(a) and (b) show the distribution of maximum IVC diameters and caval indices observed by the trained physicians in this study. The maximum diameters tend to fit a normal distribution centered around a mean of 17.9 mm with a standard deviation of 4.12 mm. The caval indices as measured by the physicians varied widely from 0-100%, with a predominant peak between 10-30%. Thirty-two subjects (40 images) had an IVC collapse of less than 50% and 15 subjects (17 images) had a collapse of greater than 50%.

A plot of the means of and differences between the automated and physician measurements of the caval indices can be seen in Figure 2.33. The difference was found by subtracting the automated measurement from the physician measurement. Three outliers were observed at higher levels of collapse indicating certain limitations of the automated tech-




Figure 2.31: Representative example of the inferior vena cava diameter tracked over time through ultrasound and the iterative KLT algorithm. Note that the raw signal (top) is unfiltered and represents the motion of the inferior vena cava precisely. Removing the respiratory signal produces a cardiac only signal (bottom) whose amplitude varies with respiration—increasing during inspiration and decreasing with expiration.




Figure 2.32: A histogram of (a) the maximum diameters of the IVCs as measured by the clinicians and (b) the caval index of the IVCs as measured by the clinicians. While a fairly normal distribution was found for maximum diameters (centered around 18 mm), the caval indices demonstrated a skewed distribution with a long tail tending toward 100% collapse.




Figure 2.33: Bland-Altman plot (mean of measurements vs. the difference in measurements) of the caval indices as measured by physicians and the algorithm. The three white points at high levels of collapse represent outliers in the data set.

nique at higher caval indices. Excluding these outliers, the average measured difference between the two techniques is 2.12% with a standard deviation of 4.48%. More broadly, there is good agreement between the two types of measurement to within 10% for over 95% of the data.

Table 1 lists the averages, standard deviations, and ranges for the maximum IVC diameter, minimum IVC diameter, and caval indices from both the physician and algorithmic measurements.³¹ Intraclass correlation coefficients of the form ICC(2,1) from Shrout and Fleiss [161] were calculated for the maximum diameter, minimum diameter, and the caval index yielding values of 0.897, 0.967, and 0.975, respectively. All measured data, including the three identified outliers, were used to populate this table.

2.4.2.4 Discussion

This algorithmic approach to measuring the IVC diameter accurately replicates the maximum, minimum, and collapsibility to within 10% of physician reported values in over

³¹These results were remarkably in line with those found in Blehar et al. [160]

Table 2.2: The averages, standard deviations, ranges, and intraclass correlation coefficients for the maximum IVC diameter, the minimum IVC diameter, and the caval index as measured by physicians and the proposed algorithm.

	Max Diam. [mm]	Min Diam. [mm]	Caval Index [%]
Mean (physician)	17.9	12.1	36.4
Mean (algorithm)	18.6	13.1	33.5
SD (physician)	4.12	6.67	29.2
SD (algorithm)	2.88	6.85	26.0
Range (physician)	8.71 27.4	0 24.2	0 100
Range (algorithm)	9.22 29.5	1.62 26.7	2.02 86.1
ICC	0.897	0.967	0.975
ICC [95% C.I.]	[0.817 0.941]	[0.881 0.986]	[0.942 0.988]

95% of measured cases, with intraclass correlation coefficients for each of the parameters equal to 0.897, 0.967, and 0.975, respectively. Such a high degree of agreement shows that our feature tracking algorithm produces measurements functionally equivalent to trained physicians. The vast majority of the differences between the measurements from a physician and from the algorithm are less than 1 mm, stemming from a slightly different set of points measured along or at the edge of the vessel. This degree of accuracy and precision coupled with the ability to give high quality continuous tracings of the IVC provides strong evidence for utilizing this technique for future work.

However, limitations to the proposed algorithm exist. The feature to be tracked, the edge of the vessel, must be clear and consistent. The algorithm as developed and reported has no means by which to correct for cylinder tangents, highly noisy images, or ultrasonic shadows, as may appear if imaging between ribs, falling along the edge. That is, the featuring tracking portion of the algorithm cannot compensate for any irregularities appearing within the images from shifting imaging probes. Another possible source of error in the measurement comes from very large, very sudden movements as may accompany a forceful sniff. Though pyramidal segmentation has been employed to combat these effects, caution is advised before implementing this algorithm to investigate very large and rapidly occurring IVC collapse.

2.4.2.5 Future work

Of particular interest is the potentially physiological meaning of both the CCI and CCIV. These may represent new diagnostic and therapeutic targets. However, it will be necessary to clearly track these in real time to better understand their importance and to this end, the

algorithm described in this report may help. I am currently investigating these parameters and their change in response to both positive and negative fluid challenges in spontaneous and mechanically ventilated patients. It is only due to time constraints that the details of that investigation are not included here.

2.4.3 Conclusions

Herein I have demonstrated the feasibility of an image processing tool to continuously measure physical parameters of the vasculature including the cardiac and respiratory phasicity of the IVC diameter as accurately and more rapidly than a trained physician, the small pulsatile motions of arteriovenous fistula as they develop over time, the elastic moduli of surface level vessels, reconstruct the three-dimensional structure of blood vessels from handheld ultrasound probes alone, and in some subset of cases even obtain an approximate measure of the shear rate of blood against the wall. These methods are combined into an open-source toolbox, providing clinicians and researchers access to underlying physiology not immediately available through previous imaging techniques. In the broadest sense, I hope that these tools encourage more to adopt advanced image processing techniques in general (and my specific open-source toolbox, specifically) to aid in their clinical observations—helping guide diagnostic decisions, helping treat patients.

Thus validated, these ultrasound techniques serve as the standard against which the bioimpedance techniques will be compared.

CHAPTER 3

Bioimpedance monitoring

And wonders within there yet.

Walt Whitman, *I Sing the Body Electric*

3.1 Theoretical basis of bioimpedance

Bioimpedance is the historically encumbered term¹ used to describe the study of tissues subject to external electrical excitation. It is distinct within the field of electrophysiology in that it is primarily concerned with ascertaining the *passive* electrical properties of biological tissues (whether living, dying, or dead) by applying a stimulus to a material, measuring the response, and inferring something of the underlying electrical, chemical, anatomical, or physiological mechanisms. This approach can be seen as distinct from the study of bioelectricity, which concerns itself with the *active* electrical properties of tissues. In this context active properties include those that generate electrical phenomenon, such as the neurons of the brain or the sinoatrial nerves of the heart. Passive properties, in contrast, are those that arise from an applied electrical stimulus.

Biological tissues, in addition to being great and wonder filled things, are also at root complicated bags of salty water. Boundaries are formed between inner cellular components and extracellular components by semipermeable membranes, a vast array of ions, organic compounds, organelles, and even other living organisms may all exist within the bounds of a particular biological tissue. This muddles matters immensely as it means tissues act as both conductors—materials that allow the flow of electrical current—and dielectrics—materials that do not allow current flow, but can be polarized in the presence of an electric field. This means that there are some facets of tissues that allow current to flow and others

¹I say historically encumbered because it has become entrenched through repetition rather than reason. The field of bioimpedance and those working within it have broader interests than the "impedance" of the name would imply. Indeed, even the "bio" prefix is not necessarily applicable to all work within the field.

which oppose it (through energy storage, not just dissipation). Further extenuating the situation is the fact that at sufficiently high frequency alternating currents, everything acts like a dielectric material, indicating that any categorical distinctions to be made are frequency dependent. Indeed, this can be described more robustly by defining a material whose capacitive current is larger than its in-phase current, such that $\omega C > G$, where ω is the angular frequency, C is the capacitance, and G is the conductance.² Thus whether one regards a tissue as mostly a conductor or mostly a dielectric will depend on the tissue itself and the context in which it is being evaluated.³

3.1.1 Permittivity and polarization

To understand what electricity does to a material, let us first review a few fundamentals. The first to concern us is Coulomb's law specifying the static relationship between force and electrically charged particles to be inversely proportional to the square of the distance between them (seen in Figure 3.1). Unlike other Mathematically this is

$$|\mathbf{F}| = k_e \frac{q_1 q_2}{r^2} \quad (3.1)$$

where $|\mathbf{F}|$ is the magnitude of the force vector, k_e is Coulomb's constant ($k_e = 8.987 \cdot 10^9 \text{ N} \cdot \text{m}^2/\text{C}^2$), q_1 and q_2 are the two electric charges, and r is the distance between the charges. Though it might not be readily apparent, Coulomb's constant actually hides within itself a very important physical parameter: permittivity.

Defined as the resistance encountered when forming an electric field within a medium, permittivity describes how a dielectric material is affected by, and itself affects, the establishment of an electric field.⁴ Coulomb's constant is related to the permittivity of free space (the permittivity of a vacuum), ϵ_0 , by

$$k_e = \frac{1}{4\pi\epsilon_0} \quad (3.2)$$

with meaningful units of newtons per coulomb, volts per meter, and tesla meters per sec-




Figure 3.1: Point charges demonstrating Coulomb's law.

²From discussion that follows, one can see that this is equivalent to $f > \sigma/2\pi\epsilon$.

³For instance, skin tends to act more as a dielectric than the more conductor-like muscle tissue that lies below it. How one models a limb, therefore, is a matter of preference as the mathematics ought to come out the same in either case.

⁴Please note another unfortunate word choice in the field here as permittivity is a measure of resistance, where as its root word "permit" suggests the opposite. Adding further to our grief is the fact that the term resistivity, discussed further on, enjoys popular parlance to refer to something else entirely.

ond. This vacuum permittivity can also be thought of as the absolute permittivity against which all other materials' permittivities can be scaled. In this way we can define the actual permittivity of a material, ϵ as the product of relative and absolute permittivity, ϵ_r and ϵ_0 , respectively:

$$\epsilon = \epsilon_r \epsilon_0 \quad (3.3)$$

Equation 3.1 can be represented vectorially by defining a vector pointing from charge one to charge two, $\mathbf{r}_{21} = \mathbf{r}_1 - \mathbf{r}_2$, and by noting the unit vector becomes $\hat{\mathbf{r}}_{21} = \mathbf{r}_{21}/|\mathbf{r}_{21}|$. Further substituting the results of equations 3.2 and 3.3 into 3.1 we are left with a medium specific Coulomb relationship for two particles

$$\mathbf{F}_1 = \frac{q_1 q_2}{4\pi \epsilon_r \epsilon_0 |\mathbf{r}_{21}|^2} \hat{\mathbf{r}}_{21} \quad (3.4)$$

One interesting result of the above equation is that larger values of permittivity produce smaller forces. This helps to explain how water with its high relative permittivity is able to dissociate many solutes that are otherwise held together by Coulomb forces.

Another way to characterize pairs of electric charge is to think of them as dipoles, doublets of positive and negative electric charges kept a distance away from each other. Such a phenomenon is described by its dipole moment, \mathbf{p} . In the simple case of two point charges separated by a displacement vector, \mathbf{d} , and with each charge, q_1 and q_2 , having the same magnitude, q , and opposite signs, the dipole moment is written as

$$\mathbf{p} = q\mathbf{d} \quad (3.5)$$

We are justified in easily extending this formulation to many such dipoles existing within a medium as

$$\sum_{i=1}^n \mathbf{p} = \sum_{i=1}^n q_i \mathbf{d}_i \quad (3.6)$$

Coulomb's law⁵ can be further expanded to give rise to the concept of the electric field, wherein a single charge at a particular point experiences a force gradient of the form

$$\mathbf{E} = \frac{\mathbf{F}}{q} \rightarrow |\mathbf{E}| = \frac{|q|}{4\pi \epsilon_r \epsilon_0 r^2} \quad (3.7)$$

The energy required to a move a charge, q , from infinity to some point across this field is $q\Phi$, where Φ is the potential at the location of interest. This potential, often called

⁵The resulting potential of dipoles within a field is extensible by the same logic.

voltage⁶, is the work done by the electric field to carry move the charge from infinity to point. Differences in electrical potential represent the amount of work that could be done if electric current was allowed to flow. Indeed, much like fluid flow's relationship to pressure, current must flow down potential gradients. Hence, the energy of the field, \mathbf{E} is equivalent to the gradient of the potential in space:

$$\mathbf{E} = -\nabla\Phi \quad (3.8)$$

Note that electric potential, Φ , is a scalar component, not a vector. The directionality of energy follows from the gradient vector, ∇ , of potential, not from potential itself.

From this cursory examination we have established how energy and the electric field are related and how current (that is, charge) flows in the presence of such a field. This latter phenomenon, namely the disturbance of the distribution of charge due to an electric field, will be referred to here as polarization and it is a key concept in the evaluation of electrical properties. Unless otherwise specified, I will use this definition when referring to polarization within a material. Within tissues, this polarization may either by endogenic (produced by a tissue itself) or exogenic (caused by an externally applied field). While the membranes of living cells are endogenic polarized (due to the sodium-potassium pump transferring ions across the cell to create a negatively charged interior with respect to the surrounding environment), the techniques of bioimpedance are almost wholly concerned with exogenic polarization.⁷

In general there are four main types of dielectric polarization: ionic, dipolar, atomic, and electronic. Ionic polarization emanates from the displacement of ions, atoms or molecules with an imbalance of charged particles, within a medium. Dipolar polarization, also called orientational polarization, results from the realigning of dipoles, either permanent or induced, when subjected to an electric field. Atomic polarization occurs when the nucleus of an atom reorients itself in response to an external electric field. Electronic polarization happens when an electric field induces very small translational displacements in the electron cloud surrounding an atom. Each of these types of polarization happens over an increasingly short period of time, respectively, with ionic polarization occurring over the longest period of time (just about everything greater than one microsecond), followed by dipolar, atomic, and finally electronic polarization (transpiring over picoseconds). Thus, depending on the frequencies present in the variation or establishment of an electric field will act by different mechanisms to polarize the dielectric.

⁶The unit for potential, the volt (equal to joules per coulomb), named for the pioneering scientist of electricity Alessandro Volta, bears some responsibility in this more conventional, if less helpful, name.

⁷Recall that bioimpedance is the study of passive electrical properties, not active ones.

Polarization is the cumulation of all dipole moments within a volume. Given the variety of different polarization mechanisms, this resultant can include dipoles from atoms, molecules, and collections of molecules. As such, polarization is the macroscopic representation of the net summation of the dipoles within a region, and its mathematical representation, \mathbf{P} is a vector in space and its orientation is a function of both energy field and another field known as the electric displacement field. The electric displacement field, \mathbf{D} , also known as the electric dipole moment density, specifies the effects of free and bound charges within a material with respect to electromagnetic phenomenon. Phenomenologically, \mathbf{D} is defined as

$$\mathbf{D} := \epsilon_0 + \mathbf{P} \quad (3.9)$$

In linear, homogeneous, isotropic dielectric materials with a negligible response time over the period of measurement, the polarization vector is equivalent to the product of the electric susceptibility and the energy field such that

$$\mathbf{P} = \epsilon_0 \chi \mathbf{E} \quad (3.10)$$

where χ is the electric susceptibility of a material ($\chi = \epsilon_r - 1$). Thus, substituting the above equation into equation 3.9 the surface charge density (another way of stating the electric displacement field) of such a material would be

$$\mathbf{D} = \epsilon_0(1 + \chi)\mathbf{E} = \epsilon\mathbf{E} \quad (3.11)$$

with units of C/m^2 for a surface charge density and Cm/m^2 for a dipole moment volume density. Both interpretations are widely used.⁸

Though I have gone one at some length, one important fact to note is that polarization cannot actually be measured. What can be measured are certain dielectric properties (resistance, capacitance, etc.) and from that polarization can be inferred and calculated. This is why we will primarily concern ourselves with the physical parameters of polarization (such as impedance) rather than its actual manifestation. Whether such a refinement (from properties to phenomenon) would improve the work demonstrated here remains to be seen. For now, we will have to proceed with the empirically based assertion that the physical properties we deal with have some bearing on the electrical phenomenon we exploit.

⁸Note further that permittivity is by this reckoning a measure of the dipole moment density induced by an electric field.

3.1.2 Complex permittivity

From equation 3.11 we can see that permittivity directly accounts for the dipole moment density induced by an electric field. By measuring permittivity (or the related parameters of conductivity or resistivity), one can infer polarization and from that back out the physical changes occurring during the measurement. Researchers in the bioimpedance field have devised a number of interesting ways to infer phenomenon including the behavior of cells (Coulter counters, hematocrit measurement, and cell culture development), the classification of tissues (detecting fluid compartment, body mass indices, and pathological tissues), and, most relevant to this work, volume changes (cardiography, pneumography, and plethysmography).

At this point it may be useful to think of permittivity as a complex quantity having both real and imaginary components. We can thus define this complex permittivity, ϵ , by convention as

$$\epsilon = \epsilon' + j\epsilon'' = (\epsilon'_r + j\epsilon''_r) \epsilon_0 \quad (3.12)$$

where ϵ' is the magnitude of the real portion and ϵ'' is the magnitude of the imaginary component. One may also wish to apply this logic to both conductivity, the measure of a material's ability to conduct an electric current (with units σ/m), and resistivity, a measure quantifying how strongly a material opposes the flow of electric current. These two quantities can be defined, respectively, as

$$\sigma = \sigma' + j\sigma'' \quad (3.13)$$

$$\rho = \rho' + j\rho'' \quad (3.14)$$

analogously to permittivity. Note that while the relationship between resistivity and conductivity is commonly written as $\rho = 1/\sigma$, what is meant is $|\rho| = 1/|\sigma|$.⁹

With these definitions we can consider a simple dielectric material sandwiched between two electrodes subjected to an alternating source (as seen in Figure 3.2). If we simplify this situation by specifying there is no direct current conductance, no accumulation of free




Figure 3.2: Simple capacitor representation of dielectric material.

⁹When dealing with only real values this does not pose much of a concern, but when working with complex signals as we do here we must bear in mind that $\rho' \neq 1/\sigma'$.

charges at the interfaces, and no electrode polarization, we are left with a picture of an ideal capacitor. The imposed electric field disturbs the distribution of bound charges in the dielectric material. Such a material is defined by $C = Q/V$, where C is the capacitance of the material, Q is the charge, and V is the potential across the material. Assuming a unidirectional, constant electric field, the magnitude of the potential across the plates separated by a distance, d , can then be calculated by the line integral (along the direction of application, x) of the electric field, $|\mathbf{E}|$ such that

$$\Phi = \int_0^d |\mathbf{E}| dx \quad (3.15)$$

Assuming a uniform charge density $\rho_q = Q/A$, where Q is the charge over the electrode area, A , and noting that $|\mathbf{E}| = \rho_q/\epsilon$, the above equation can be solved as

$$V = \frac{Qd}{\epsilon A} \quad (3.16)$$

Noting further that the capacitance in such a scenario is equal to

$$C = \frac{A}{d} \epsilon \quad (3.17)$$

However, once one applies an alternating signal (be it potential or current) through the electrodes into such a material, it will become immediately apparent (with the right tools) that the maximum current and maximum potential will not occur at the same time. There is a delay in time between the two that we call a phase-shift and in an ideal, steady-state capacitor subjected to a sinusoidal signal, the voltage will lag behind the current by 90. Put another way, the conductive component of the signal and the capacitive component of the signal are separated by a phase of 90. As per convention we call the conductive component *in-phase* and the capacitive component *out-of-phase*. Unlike DC signals, then, it is not possible to characterize an AC signal or response with a single number and minimum of two are required. These two numbers describe the relative proportion of an AC signal that is in-phase and out-of-phase.

In the capacitor model, therefore the total admittance—the measure of how easily current flows through a material—has both an in-phase component, conductance, G , and 90 phase-shifted component, susceptance, B . Admittance, \mathbf{Y} , is thus written as

$$\mathbf{Y} = G + jB \quad (3.18)$$

and in the case of the dielectric material we have described (substituting equations 3.12,

3.30, and 3.17) can be written as

$$\mathbf{Y} = \frac{A}{d} (\sigma' + j\omega\epsilon') \quad (3.19)$$

Since conductivity and permittivity are related by

$$\sigma = j\omega\epsilon \quad (3.20)$$

a wide range of additional relationships between the passive electrical parameters can be written. While I encourage readers to look into this matter further, for now, those relationships will be assumed in what follows and used as needed.

3.1.3 Relaxation and dispersion




Figure 3.3: Dielectric responses. Modified from "Dielectric responses" by Archimerged found at Wikipedia. Licensed under Attribution via Commons - <https://commons.wikimedia.org/wiki/>

bye in 1913. This model (the Debye single dispersion model), assumes a non-interacting population of dipoles in the presence of an alternating exogenic electric field, posits

$$\epsilon(\omega) = \epsilon'_{\infty} + \frac{\epsilon'_s - \epsilon'_{\infty}}{1 + j\omega\tau} \quad (3.21)$$

As we have mentioned previously, there is a frequency dependence to permittivity, a property known as dispersion, which is caused by the lag between the application of an electric field and the polarization that results. This lag in the system is known as dielectric relaxation and is a function of both the material and the applied field's frequency. Dispersion and relaxation ultimately describe the same phenomenon, with the former doing so in the frequency domain and the latter in the time domain.

One of the most widely used models of dielectric relaxation was first introduced by Peter De-

where $\epsilon(\omega)$ is the complex permittivity varying with angular frequency, ω , ϵ'_∞ is the permittivity at an infinitely high frequency, ϵ'_s is the static, low frequency permittivity, and τ is the characteristic relaxation time constant of the medium. It is quite common for materials to have different relaxation time constants and thus different dispersions across a broad range of frequencies. As stated before, different polarization mechanisms take different periods of time to happen over, and as such are limited in their frequency responses: dipolar polarization cannot respond to oscillations greater than about 10 GHz (in the microwave region), ionic polarizations are generally limited to electric fields below 10 THz (within the infrared region), and electronic polarization loses its ability to respond above 1 PHz (just above the ultraviolet region). Therefore, many distinct regions can characterize a dielectric material's dispersion as a function of frequency. One general example of such dispersive behavior for complex permittivity can be seen in Figure 3.3 where regions are clearly demarcated by ionic, dipolar, atomic, and electronic polarization.

3.1.4 Biological materials and their electrical equivalents

While the Debye model of relaxation works well for many materials, researchers in the 1920-1930s discovered that it did not quite match the results found when testing biological materials. Because even the simplest biological materials have immense variety in their material composition, the idea of a distribution of relaxation time constants was proposed to compensate for the differences in empirical observations and theoretical expectations. Early work by Kenneth Cole [162, 163, 164, 165, 166] eventually led to a model of dielectric materials (developed with his brother Robert Cole) that extended the logic of the Debye model by introducing the concept of a constant phase element—a component whose phase angle is independent of frequency and thus stores energy in addition to dissipating it. This model, known as the Cole-Cole model, was first introduced in 1941 [167, 168], adds a parameter, α , to the exponent of the $j\omega\tau$ term in equation 3.21 that defines the distribution of relaxation times in a material:

$$\epsilon(\omega) = \epsilon'_\infty + \frac{\epsilon'_s - \epsilon'_\infty}{1 + (j\omega\tau)^{1-\alpha}} \quad (3.22)$$

This is the Cole-Cole model so revered in the bioimpedance community. Its simplicity belies its ingenuity: with this single additional parameter, the vast majority of biological tissues can be described accurately, and has been done so for decades. However, as Cole and Cole noted the physical significance of this parameter is unclear [167] and the distribution of relaxation times it proposes is not immediately or intuitively obvious. Indeed many different distributions have been observed in dielectric materials and their interpretation is




Figure 3.4: Idealized behavior of dielectric dispersions for biomaterials showing α , β , and γ regions. Modified for educational purposes from *Bioimpedance and Bioelectricity Basics* by Grimnes and Martinsen, originally from Schwan.

still a matter of debate.¹⁰

In biological materials generally, the dielectric spectroscopic response curves tend to display three general regions of dispersion— α -, β -, and γ -dispersion—as seen in Figure 3.4.¹¹ These dispersions arise from frequency- and relaxation-dependent mechanisms: *alpha*-dispersion is believed to be caused by active cell membrane and gated channel effects, ionic diffusion, and intracellular structures; *beta*-dispersion reflects passive membrane capacitance, Maxwell-Wagner effects, and the response of protein molecules; and *gamma*-dispersion results from dipolar mechanisms present in polar media (water, salts, etc.).

The presence of distinct regions of dispersive behavior suggest to many researchers the presence of multiple distinct bands of relaxation time constant distributions and it is quite common to see biological materials modeled as such if one measures over a sufficiently large frequency spectrum. One of the most impressive procedures to document the permittivity of a wide variety of biological tissues over an immense frequency range (from 10 Hz

¹⁰Since the distribution of relaxation times does not bear on this work any more than is needed for a recognition of its presence, we will not concern ourselves with it further.

¹¹Another type, δ -dispersion, often appears in the GHz range. As that is far outside our scope of inquiry, it will not concern us here.

to 20 GHz) was undertaken by Gabriel et al. [169, 170, 171]. In this work, the permittivity of nearly 50 biological tissues was surveyed from the literature, measured, and independently fit to a parameterized summation of four Cole-Cole models.¹² That model is of the form

$$\epsilon(\omega) = \epsilon_{\infty} + \sum_{i=1}^4 \frac{\Delta\epsilon_i}{1 + (j\omega\tau_i)^{(1-\alpha_i)}} + \frac{\sigma_s}{j\omega\epsilon_0} \quad (3.23)$$

where i is the indexing variable for each Cole-Cole model, $\Delta\epsilon$ is the difference between the low and high frequency permittivity ($\epsilon_s - \epsilon_{\infty}$), and σ_s is the static ionic conductivity. Table 3.1 is populated with the values fit for each of the parameters for a few biological materials of interest.

To simplify matters, rather than model tissues as dielectrics, biological materials may also thought of in terms of equivalent impedance of a simple circuit.

Electrical impedance is a measure a materials resistance to an electrical flow given a potential. This can be understood through Ohms law, which states

$$\mathbf{Z} = \frac{\mathbf{V}}{\mathbf{I}} \quad (3.24)$$

where \mathbf{Z} represents the impedance vector, \mathbf{V} represents the voltage vector and \mathbf{I} represents the current vector. Because each of these variables is complex, they have real and imaginary components. The two components that make up the impedance vectors are the resistance, R , and the reactance, X , constituting the magnitudes of the real and imaginary elements of \mathbf{Z} . Mathematically, this can be written as

$$\mathbf{Z} = R + jX \quad (3.25)$$

where the reactance is augmented by the imaginary unit equal to the square root of -1 in much the same way equation 3.18 is. (Indeed impedance is merely the inverse of admittance, $\mathbf{Z} = 1/\mathbf{Y}$) From basic vector mathematics, we know that the magnitude of impedance is

$$|\mathbf{Z}| = \sqrt{R^2 + X^2} \quad (3.26)$$

and the phase is

$$\phi = \tan^{-1} \left(\frac{X}{R} \right) \quad (3.27)$$

For most tissues, the magnitude and phase of impedance are frequency dependent (an

¹²As the authors acknowledge, though only three main dispersions are evident over the frequency range they measured, a four term model was used for flexibility and a better fit.

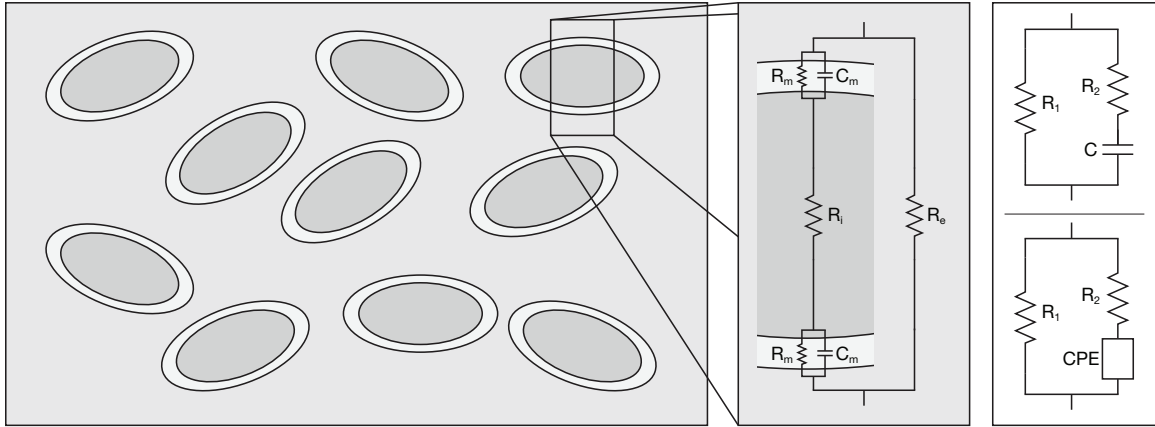


Figure 3.5: The cells of a tissue and their electrical component equivalents. For an individual cell, R_m and C_m are the resistance and capacitance of the cellular membrane, R_i is the intracellular resistance, and R_e is the extracellular resistance. Summing up all the cells yields one of two popular equivalent circuits to work with, one (right top) the parallel combination of a resistor, R_1 with a series combination of a resistor, R_2 with a capacitor, C , and the other (right bottom) the parallel combination of a resistor, R_1 with a series combination of a resistor, R_2 with a constant phase element, CPE .

by extension, so too are resistance and reactance are). This is due to the combination of cellular (intracellular elements surrounded by cellular membranes) and extracellular components. The cellular membranes separating the resistive intra- and extracellular fluids act much like capacitors or constant phase elements, themselves a time/frequency dependent electrical component that behave as high pass filters (low frequency alternating currents do not pass through easily, high frequency do). Such electrical behaviors have led research to treat biological tissues as an equivalent circuit composed resistor in parallel with a series combination of a resistor or a capacitor (as shown in Figure 3.5).

From this equivalent circuit and from observation, researchers know that if such a system were subjected to a frequency sweep, the tip of the bioimpedance vector traces out a depressed semi-circular locus on a resistance-reactance plot as shown in Figure 3.6. The path traced is a function of the intracellular and extracellular fluids within the body, themselves a function of the sum total of the volumes of conductors and nonconductors. Put simply, so modeling impedance gives investigators another means by which to represent the dielectric properties of a material.




Figure 3.6: Representative resistance-reactance plot for biological tissues. $\mathbf{Z}(\omega)$, $\theta(\omega)$, $R(\omega)$, and $X(\omega)$ are the impedance, phase angle, resistance, and reactance at an angular frequency, ω , R_0 is the resistance at low frequencies, R_∞ is the resistance at high frequencies, X_{peak} is the peak reactance and it occurs vertically over the depressed locus, D .

Table 3.1: Dielectric properties of a range of biological tissues. The parameters fit to the parameterized Cole-Cole summation as seen in equation 3.23. Here, k and M represent multiplying the value by 10^3 and 10^6 , respectively. The materials displayed here are the same as those reported in 2.1.

Material	ϵ_∞	$\Delta\epsilon_1$	τ_1 [ps]	α_1	$\Delta\epsilon_2$	τ_2 [ns]	α_2	$\Delta\epsilon_3$	τ_3 [μ s]	α_3	$\Delta\epsilon_3$	τ_4 [ms]	α_4	σ_s
Air	1	0	0	0	0	0	0	0	0	0	0	0	0	0
Blood	4	56	8.4	0.1	5.2k	132	0.1	0	159	0.2	0	15.9	0	0.7
Blood Vessel Wall	4	40	8.8	0.1	50	3.2	0.1	100k	159	0.2	10M	1.6	0	0.25
Bone (Cortical)	2.5	10	13.3	0.2	180	79.6	0.2	5k	159	0.2	100k	15.9	0	0.02
Brain	4	40	8.0	0.1	700	15.9	0.15	200k	106	0.22	45M	5.3	0	0.04
Fat	2.5	9	8.0	0.2	35	15.9	0.1	33k	159	0.05	10M	15.9	0.01	0.035
Heart Muscle	4	50	8.0	0.1	1.2k	159	0.05	450k	72.3	0.22	25M	4.5	0	0.05
Kidney	4	47	8.0	0.1	3.5k	199	0.22	250k	79.6	0.22	30M	4.5	0	0.05
Liver	4	39	8.8	0.1	6.0k	531	0.2	50k	22.7	0.2	30M	15.9	0.05	0.02
Lung	2.5	18	8.0	0.1	500	63.7	0.1	250k	159	0.2	40M	8.0	0	0.03
Mucous Membrane	4	50	7.2	0.1	7.0k	354	0.1	1.2M	318	0.1	25M	2.3	0	0.2
Muscle	4	50	7.2	0.1	7.0k	354	0.1	1.2M	318	0.1	25M	2.3	0	0.2
Nerve	4	26	8.0	0.1	500	106	0.15	70k	15.9	0.2	40M	15.9	0	0.006
Skin	4	32	7.2	0	1.1k	32.5	0.2	0	159	0.2	0	15.9	0.2	0.0002
Urine	3	66	6.9	1	28	1.8	1	0	159	0	0	15.9	0	1.75
Water	0	0	0	0	0	0	0	0	0	0	0	0	0	0

3.2 Practical application of bioimpedance

3.2.1 Clinical use

Many techniques use the passive electrical properties of materials to glean information about them. Bioimpedance specifically has been used in non-medical applications as diverse as authenticating finger prints [172], tracking the fermentation of alcohol [173, 174], and monitoring the quality of meat [175, 176]. Researchers working with everything from single cells [177] and cultures [178] to tissues [179] and organs [180] to whole bodies [181] and populations [182] have used bioimpedance based methods to quantify anatomy and physiology, structure and function.

As might be expected, a significant portion of bioimpedance researchers have focused their efforts on medical applications, where the dielectric properties are reflective of underlying chemistry and biology. In this section, I will briefly review three pillars of bioimpedance relevant to the research presented in the latter half of this chapter¹³: static, dynamic, and multidimensional approaches.

In this context, both *static* and *dynamic* approaches refer to techniques yielding a single value at a given time for their description. That is, the technique yields a single number (and it may be a complex number) and can thus be thought of here as “one-dimensional.” Whether this value is expected to remain constant or vary over the course of measurement distinguishes the static from the dynamic, and examples of each can be seen in Sections 3.2.1.1 and 3.2.1.2, respectively. In contrast, a *multidimensional* approach, such as electrical impedance tomography, synthesizes many one-dimensional measures to give a fuller picture of the observed material. Such an approach may itself be either static, dynamic, or some combination of the two as is the case in the example given in Section 3.2.1.3.

Though such a parsing is convenient for this discussion, one must not act as though these distinctions are hard and fast. For instance, one could use a static measurement over multiple time frames, along multiple sites, or across multiple people within a population and my delineation of it here as unchanging is no longer as neatly drawn. Furthermore, dynamic measures need not vary as a function of time. Mechanical deformation [183] and biological development can be readily described in terms of bioimpedance. This framework was chosen to better triangulate the place of my bioimpedance technique and the examples elucidated below are presented to the same ends.

¹³I leave other equally practical techniques applying bioimpedance such as for the control of electrosurgery, the measurement of electrodermal activity, and rheoencephalography for the interested reader to discover.

3.2.1.1 Bioimpedance analysis: fluid distribution and hydration status

One of the most common and widely implemented uses of “static” bioimpedance techniques clinically is the measurement of fluids throughout the body [184]. Intracellular fluids, extracellular fluids, their summation (known as total body water content, whole body fluid content, etc.), and their distribution throughout the body all factor critically into normal physiological functioning and as such have been used to diagnose the extent and severity of many pathologies.¹⁴

Broadly speaking, there are four general bioimpedance based methods to assess fluid volumes throughout the body.

1. **Single frequency bioimpedance analysis.** The most reported method is the prediction of total body water from whole body impedance measurements, typically taken at a single frequency. To compensate for individual variations, several other variables are often introduced—from gender to age, from height to weight—in the desire to increase accuracy. Many equations have been found and experiments conducted demonstrating good empirical correlation between body and segmental impedance measures and total body water [181, 185, 184, 186, 187], though such relationships have not thus far been directly derived from biophysical models.
2. **Bioimpedance spectroscopy.** A second uses bioimpedance spectroscopy to fit a model (recall 3.6), often of the parallel branch.¹⁵ The resistance values at either extreme, R_0 and R_∞ , are extrapolated to predict the extracellular and total body fluid components, respectively. The intracellular fluid components can be found by taking the difference between these two measures. By assuming certain facts about the fluids’ resistivities, the volumes of each components can be estimated. The anisotropy of real tissues and the range of resistance values possible from the typical to the pathological conspire to make this technique useable in only narrow and specialized circumstances. Evidence suggests it may be useful in monitoring during dialysis [188].
3. **Multiple frequency bioimpedance analysis.** A combination of the first two, a third method estimates extracellular and total body water content by inputting only low (1-5 kHz) and high frequency (>100 kHz) signals, respectively. Intracellular water

¹⁴The extracellular fluid compartment is composed of two distinct spaces: the intravascular and the interstitial. I relegate this to a footnote for now, but its distinction is crucial for the proposed use (outlined in Section 3.2.2).

¹⁵As an aside, throughout the literature this is all too often referred to as a “Cole-Cole” model. This is wrong. The Cole-Cole model, you will recall, is a permittivity model and does not speak to resistance or reactance in this way. I applaud the efforts of Grimnes and Martinsen to dissuade this sort of terminology.

is still found by taking the difference between the two and requiring a slate of assumptions: tissue isotropy, low frequency current not penetrating cells, fixed fat free mass (almost always assumed to be 73% in these cases), etc. In cases where such assumptions are not valid (during hemodialysis, patients with edema, those with heart failure, pregnant women, athletes, etc.), such techniques ought not be employed.

4. **Bioimpedance vector analysis.** Referenced early in Section 1.6.4.4, bioimpedance vector analysis uses both the real and imaginary components of the impedance vector on a resistance-reactance diagram to probabilistically determine hydration and body composition status from impedance normalized to the conductor length (height in the case of whole body, the length of a segment in segmental cases). Often a single frequency is employed (usually 50 kHz) and hydration status and body composition is determined by where the vector lies within elliptical (sex- or condition-specific) probability regions reported in the literature. A sizable body of evidence suggests this approach is useful to compare an individual against the greater population or against themselves across time [189, 190, 191] for these parameters.

Though it was introduced half a century ago [192] and has seen near continual refinement since then, total body impedance to assess fluid volume content remains controversial. For one, the limbs and digits dominate the measure since their small cross-sectional areas increase resistance regardless of material composition. Further complicating the thoracic contribution are the continual changes due to respiration and the cardiac cycle. Though segmental approaches—wherein portions of the body are measured individually and are either compared or summed—attempt to diminish these issues, standardization of electrode placement and type remain concerns [187, 186].

Body position and habitus greatly influence fluid volume measurements of this kind because of the combined contributions of gravity and the paths available for fluid to travel, respectively. Anyone wishing to use bioimpedance in these




Figure 3.7: Representative bioimpedance vector analysis graph. The impedance vector normalized to the patient height falls within a probabilistic region to determine the ratio of fluid to soft tissues a person has.

ways must recognize the limitations of such an approach and compensate.

3.2.1.2 Impedance plethysmography: stroke volume and cardiac output

Impedance plethysmography—measuring volume changes via conductance changes—was famously reported in the literature by Nyboer in 1950 [193]. In this landmark study, small changes in blood volume in a limb corresponded to measurable changes in impedance, thus opening vistas of future research into peripheral and central arterial and venous blood flows to examination via bioimpedance (including my own).

Clinically this has found a foothold in impedance cardiography, a technique in which the measurement of stroke volume and therefore cardiac output can be calculated by measuring thoracic electrical bioimpedance over time. Regarding the thorax as a simple cylinder with a length, L , researchers (starting with Nyboer) describe the change in volume, Δv in terms of the change in impedance, ΔZ as

$$\Delta v = \rho \left(\frac{L}{Z_0} \right)^2 \Delta Z \quad (3.28)$$

where ρ is the resistivity and Z_0 is the initial impedance. From this stroke volume, SV was first defined by Kubicek et al. [194] as

$$SV = \rho \left(\frac{L}{Z_0} \right)^2 \left(\frac{dZ}{dt} \right)_{max} \cdot LVET \quad (3.29)$$

where $\left(\frac{dZ}{dt} \right)_{max}$ is the maximum value of the first time derivative of impedance and $LVET$ is the left ventricular ejection time (usually estimated from a correspond electrocardiographic measurement). This model corresponds well to a the basic physics of the environment and is still used widely [195, 196, 197, 198].

However, fluid within the chest moves in a more complicated manner than as a just a simple cylindrical addition. Measuring from the chest, plethysmographic signals arise from the filling and emptying of the heart, aorta, vasculature, lungs, and muscles, each of varying cross-sectional areas and constitutive materials. Moreover, the assumption that all measured variations in impedance arise from changes in an evenly distributed electrical field is not fully satisfied in a thoracic cavity with multiple organs with myriad dielectric properties (recall Table 3.1).

As I have pointed previously (in Chapter 1) and as Grimnes and Martinsen put forcefully [199], there simply is no gold standard for stroke volume and cardiac output measurements. The simple and low-cost nature of impedance cardiography as a technique appeals




Figure 3.8: An example of electrical impedance tomography demonstrating (a) healthy and (b) unhealthy lungs. The pixel intensities here represent a range-scaled map of electrical impedance measured across sixteen electrodes along the subject’s chest (represented by the perimeter of white dots in each image). In the case of the healthy lungs, air (and thus impedance) is evenly distributed between the front and back of both lungs. The unhealthy lungs, from a subject under mechanical ventilation undergoing tidal recruitment, show one lung and one portion of that lung preferentially letting in air. The images of this figure were generated using data supplied by Drager.

to many clinicians and researchers and I suspect in the future could play a larger role in a smaller context (for instance, that of wearable medical monitors).

3.2.1.3 Electrical impedance tomography: lung imaging and beyond

Electrical impedance tomography expands bioimpedance signals from their aforementioned one-dimensional representations to multidimensional images (usually two-dimensional images played in sequence, though volumetric presentations are also possible) [200]. This has found a niche in lung monitoring for patients under mechanical ventilation [201, 202] (an example of which can be seen in Figure 3.8). Work has also been conducted to demonstrate its applicability for continuous intracranial monitoring to track cerebral edema, stroke, or intracranial hemorrhage [203, 204, 205].

3.2.2 Proposed use

As we have seen, impedance plethysmographic techniques capable of measuring changes in blood volume have been frequently used since the 1950s. For the work presented here, a novel bioimpedance based plethysmographic technique, deemed impedimetric intravascu-

lar volume evaluation (I2VE), proposes that changes in limb impedance can be treated as a surrogate of the volume status and responsiveness of a subject. For this to be true a number of assumptions must hold. Those assumptions are as follows:

1. **Changes in conductance over the respiratory cycle are due exclusively to changes in blood volume present in the limb.** All other sources of impedance variation must be negligible (such as long term tissue development) or separable (such as muscle contractions). Were this initial assumption not satisfied—if impedance values changed significantly over the respiratory cycle with no respect to blood content in the limb—this impedance plethysmographic technique would have no bearing.
2. **The blood volume changes occurring over the respiratory cycle must be correlated with respiration.** Intrathoracic pressures must be able to cause some measurable shift in blood within the vasculature. This assumption is necessary to link impedance with cardiorespiratory physiology.
3. **The respiratory induced variations in blood volume must be measurably different.** If the changes seen are below the level of sensible difference it would not be sensible to proceed with the technique.

With these assumptions satisfied we may proceed.

While impedance plethysmography is nothing new, what I believe is unique about this particular use case is its foundation: an intrinsic link between anatomy and physiology that yields actionable information for clinical decision making. Surprisingly, perhaps dishearteningly so, this is not the case in many instances. By linking the physiology of volume status (discussed at length in the first chapter) with the measurement of dynamically varying conductivity in the limb, new and clinically relevant information emerges about the current state of a subject. Furthermore, by perturbing a subject’s cardiovascular system with a small respiratory maneuver, the responsiveness to a treatment (such as the infusion of a fluid bolus) can be immediately gauged.

3.2.2.1 An ideal environment

To model a limb’s vasculature, let us model the entire system as consisting of a single cylinder with radius r and length L separating an interior dielectric fluid with conductivity, σ , and a medium of constant conductivity that spreads infinitely in all directions. We are left with a situation, graphically represented in Figure 3.9, where the conductance, G , inside the cylinder is

$$G = \frac{\pi r^2 \sigma}{L} \quad (3.30)$$

Thereby making the volume

$$v = G\rho L^2 \quad (3.31)$$

Assuming that conductivity remains constant throughout¹⁶, the change in conductivity is therefore a reflection of the change in the volume, Δv , of the cylinder:

$$\frac{\Delta G}{G} = \frac{\Delta v}{v} \quad (3.32)$$

From equations 3.30 and 3.31 the change in volume in the cylinder simply becomes

$$\Delta v = \Delta G\rho L^2 \quad (3.33)$$

From equation 3.33 we can infer two interesting conclusions. The first is that even without prior knowledge of the original geometry of the cylinder (or even that it is cylinder at all), we can directly link changes in conductance (or impedance¹⁷) to changes in volume. The second conclusion to be drawn is that to make the measurement as sensitive as possible the length of the observed sample should be as small as possible. In practice, however, there is a balance to be struck: an L too large diminishes the effects of the conductance changes and an L too small can make it difficult to establish a sufficient current density for measurement.

The results of equations 3.28 and 3.33 indicate that a change in volume is proportional to the change in conductor present. For either to be utilized, the change in impedance must be small in comparison to the baseline, which the literature and my personal experience shows to be true in peripheral limb impedance. The results presented in the following sections proceed with this model. To account for more complicated and realistic environments where the assumption of a simple non-deforming cylinder does not hold (for instance, the fingers, hands, wrists, thorax, etc.), I have derived the relationship between the change in impedance as volume from first principles in the section that follows.

3.2.2.2 A more realistic environment

In reality, some situation are a little more complicated: there can be more than a single vessel's impedance change to measure; the volume of blood varies with the cardiac cycle,

¹⁶Not necessarily a valid assumption for blood given its flow dependent conductivity through the Sigman effect. The ramifications of this effect is discussed later.

¹⁷While measuring either impedance or conductance yields equivalent results, the mathematics makes this less than obvious. In this case $\Delta v = \left(\frac{1}{R+\Delta R} - \frac{1}{R}\right) \rho L^2$, where R is the resistance. Assuming $\Delta R \ll R$ this becomes approximately $\Delta v \approx -\Delta R\rho \left(\frac{L}{R}\right)^2$. This result is equivalent to that presented equation 3.28.




Figure 3.9: A representation of the parameters involved in tetrapolar impedance measurement. Four electrodes, A , B , C , D are placed on a volume, v , with potentials ϕ and ψ , current densities \mathbf{J}_ϕ and \mathbf{J}_ψ , and conductivities σ_ϕ and σ_ψ . A bolus of fluid travels from A to B causing a conductance change, a potential difference, and a measurable impedance variation across the electrodes.

the respiratory cycle, and with the presence of gravity; and boluses of blood travel in one direction along the arteries and in the opposite direction in the veins. Beyond that, the bioimpedance measurements taken on the limb are made with a tetrapolar arrangement similar to that represented in Figure 3.9 with two outer electrodes injecting a small amount of current and two inner electrodes sensing potential rather than an non-dimensional bipolar arrangement assumed in the ideal case. The tetrapolar case has many practical advantages, most especially the ability to minimize the effects of electrode polarization, but requires us to move beyond our simple models and turn to lead field theory to fully explain.

Let us consider an arbitrary volume conductor with two electric fields.¹⁸ The first, caused by the application of I_ϕ to the terminal pair AB , has distribution of potential $\phi(x, y, z)$, a current density $\mathbf{J}_\phi(x, y, z)$ and a conductivity of the total volume σ_ϕ . The second, caused by the application of I_ψ through the terminal pair CD , has a potential distribution $\psi(x, y, z)$, current density \mathbf{J}_ψ , and conductivity σ_ψ throughout the volume. From the divergence theorem¹⁹ we may model the current density throughout the entire volume as

$$\iiint_v \nabla \cdot (\psi \mathbf{J}_\phi) dv = \iiint_v (\psi (\nabla \cdot \mathbf{J}_\phi) + \mathbf{J}_\phi \cdot \nabla \psi) dv = \oint_s \psi \mathbf{J}_\phi \cdot d\mathbf{s} \quad (3.34)$$

and

$$\iiint_v \nabla \cdot (\phi \mathbf{J}_\psi) dv = \iiint_v (\phi (\nabla \cdot \mathbf{J}_\psi) + \mathbf{J}_\psi \cdot \nabla \phi) dv = \oint_s \phi \mathbf{J}_\psi \cdot d\mathbf{s} \quad (3.35)$$

¹⁸The basis of this formulation can be found in Geselowitz [206] and Lehr [207].

¹⁹ $\iiint_v \nabla \cdot \mathbf{A} = \oint_s \mathbf{A} \cdot d\mathbf{s}$

Since $\nabla \cdot \mathbf{J}_\phi$ and $\nabla \cdot \mathbf{J}_\psi$ are both equal to zero and the potential fields are finite within the volume, these equations reduce to

$$\iiint_v \mathbf{J}_\phi \cdot \nabla \psi dv = \oint_s \psi \mathbf{J}_\phi \cdot d\mathbf{s} \quad (3.36)$$

$$\iiint_v \mathbf{J}_\psi \cdot \nabla \phi dv = \oint_s \phi \mathbf{J}_\psi \cdot d\mathbf{s} \quad (3.37)$$

Combining equations 3.36 and 3.37 yields

$$\iiint_v (\mathbf{J}_\phi \cdot \nabla \psi - \mathbf{J}_\psi \cdot \nabla \phi) dv = \oint_s (\psi \mathbf{J}_\phi - \phi \mathbf{J}_\psi) \cdot d\mathbf{s} \quad (3.38)$$

which, by extending the logic of equations 3.8 and 3.11 results in a formulation of

$$\iiint_v (\sigma_\psi - \sigma_\phi)(\nabla_\phi \cdot \nabla_\psi) dv = \oint_s (\psi \mathbf{J}_\phi - \phi \mathbf{J}_\psi) \cdot d\mathbf{s} \quad (3.39)$$

If the individual conductivities over the period of time over which these fields are applied are equal (that is $g = g_\phi = g_\psi$) we are left with

$$\oint_s (\psi \mathbf{J}_\phi - \phi \mathbf{J}_\psi) \cdot d\mathbf{s} = 0 \quad (3.40)$$

which is a convenient result.²⁰

If we now consider a new conductivity, g' associated with the electrode pair AB with a corresponding potential of $\phi'(x, y, z)$ and current density of \mathbf{J}'_ϕ , equation 3.39 becomes

$$\iiint_v (\sigma' - \sigma)(\nabla_{\phi'} \cdot \nabla_\psi) dv = \oint_s (\psi \mathbf{J}'_\phi - \phi \mathbf{J}_\psi) \cdot d\mathbf{s} \quad (3.41)$$

Assuming the same current as the initial case, we can then derive

$$\iiint_v (\sigma' - \sigma)(\nabla_{\phi'} \cdot \nabla_\psi) dv = \oint_s (\phi - \phi') \mathbf{J}_\psi \cdot d\mathbf{s} \quad (3.42)$$

$$\iiint_v \Delta\sigma (\nabla_{\phi'} \cdot \nabla_\psi) dv = \Delta\phi_{CD} I_\psi \quad (3.43)$$

$$\iiint_v \Delta\sigma \left(\frac{\nabla_{\phi'}}{I_\phi} \cdot \frac{\nabla_\psi}{I_\psi} \right) dv = \frac{\Delta\phi_{CD}}{I_\phi} \quad (3.44)$$

$$\iiint_v \Delta\sigma \left(\frac{\nabla(\phi + \Delta\phi)}{I_\phi} \cdot \frac{\nabla_\psi}{I_\psi} \right) dv = \frac{\Delta\phi_{CD}}{I_\phi} \quad (3.45)$$

²⁰It may also conveniently be expressed as $\oint_s \psi \mathbf{J}_\phi \cdot d\mathbf{s} = \oint_s \phi \mathbf{J}_\psi \cdot d\mathbf{s}$ if the reader wishes.

and from symmetry finally yields an expression for the change in impedance of

$$\Delta Z = -\frac{\Delta\sigma}{\sigma_0^2} \iiint_v \mathbf{J}'_0 \cdot \mathbf{J}'_{\Delta\sigma} dv \quad (3.46)$$

where \mathbf{J}'_0 is the fixed lead field at an initial time point $t = 0$ and with a conductivity of σ_0 and $\mathbf{J}'_{\Delta\sigma}$ is the time varying lead field dependent on the change in conductivity, $\Delta\sigma$. From this approach, we can see that theoretically it should not matter which pairs of electrodes inject current or sense potential. Moreover, we need not know much about the initial volume or geometry of the material under examination, we merely need to know how to attribute the change in impedance to the change in conductivity and current density. Again, in reality the situation can be a little more complicated than this. Depth, frequency, and electrode geometry and distance all significantly affect current density and potential gradients in a way that is not addressed here. All the same, this examination demonstrates that changes in conductance within a volume manifest as impedance changes, providing the foundation upon which to construct the work that follows.

3.3 Impedimetric intravascular volume evaluation

The volume of and ability to circulate blood within the cardiovascular system comprise a crucial relationship to monitor for many acutely and chronically ill patients. From dehydration to renal failure, from cardiac arrest to heart failure, from sepsis to trauma the hemodynamic parameters of a patient along the progression and recovery of a disease or injury is critical for optimal diagnosis and treatment. While many hemodynamic parameters can be measured easily for many types of patients (such as arterial blood pressure and heart rate), volume status remains elusive. Moreover, if and to what extent a patient responds favorably to fluid therapy is currently difficult to gauge in clinic short of simply performing the procedure.²¹

I2VE in this context works by applying a set of electrodes (in either a bipolar or tetrapolar configuration) to a limb (either upper or lower), monitoring the impedance change over time, and tying a respiratory maneuver induced variation to a baseline value. The maneuver to cause a respiratory induced variation, also called a respiratory challenge throughout, must result in an intrathoracic pressure difference. As stated in the assumptions of the proposed use (see Section 3.2.2) the changes in limb impedance must correspond to changes in blood volume which must itself stem from a respiratory challenge which in turn tests the venous return functions of the cardiovascular system at that time. Thus by perturbing the

²¹Recall the potential limitations of a fluid challenge in *Monitoring* of Chapter 1.




Figure 3.10: A representative spontaneous breathing baseline. Here a range normalized bioimpedance waveform (solid line) demonstrates two distinct components: a higher frequency and often lower amplitude signal is superimposed upon a lower frequency corresponding to respiration (dashed line).

system, its functioning can be determined.

In spontaneously breathing patients a negative intrathoracic pressure is generated by the diaphragm of the chest pulling down during inspiration. As the diaphragm pushing upward expelling air out of the lungs, a positive intrathoracic pressure develops.²² During inspiration, the negative intrathoracic pressure gradient extends through the upper body and a corresponding negative pressure gradient is established infradiaphragmatically (in the lower half of the body). This negative pressure gradient causes more blood in the veins of the upper body (for instance, the limbs) to return to the chest cavity, ultimately emptying into the heart. The degree to which blood is pulled from the upper limbs is a function of both how forcefully a person inspires (or expires) and the current venous return capabilities, the latter of which is a direct indicator of volume status. Blood, as a comparatively good conductor within the environment of the limbs (see Table 3.1), causes a distinguishable distinguishable trend to occur as it flows distally in the arteries and proximally in the veins. All else being equal, increased blood volume within a limb causes the impedance to decrease and vice versa. Thus, during inspiration when blood is being pulled back to the center of the body, impedance increases and during expiration the opposite is true. Figure 3.10 demonstrates this via a typical impedance plethysmographic waveform. Note that two distinct components can be clearly seen: a higher frequency signal rides atop a lower frequency signal. These two signals represent cardiac and respiratory impedance variations, respectively.

For the spontaneously breathing patient populations tested here, the normal, non-perturbed breathing signal was treated as the baseline against which the respiratory induced varia-

²²In the case of mechanical ventilation, this relationship is reversed. This is discussed further in Section ??.




Figure 3.11: An example of a bioimpedance waveform collected during forced inspiration. Two normal breaths are preceded by four forced breaths (each limited to 20 cm-H₂O [2 kPa]). A distinct positive and magnified peak-to-peak amplitude difference can be seen compared to the normal inspiratory maneuvers preceding them.

tion was compared. In these populations, two forced respiratory maneuvers were applied: forced inspiration and forced expiration. Forced inspiration tasked the venous system by drawing a larger volume blood from the upper limbs and pushing a large volume of blood to the lower limbs than typical respiratory pressures would cause themselves. This in turn caused the impedance of the upper limbs to increase significantly with inspiration (as seen in Figure 3.11) and decrease significantly in the lower limbs. Forced expiration had a somewhat opposite effect: the positive pressure of the chest resisted the venous pressure gradient, either preventing blood from returning to the chest or causing less blood to do so. In both instances this results in an increase of blood volume within the limb during expiration and as such lowers the impedance (as seen in Figure 3.12). Though both types of forced maneuver test the volume status of the patient, what was often preferred in this analysis was the forced inspiratory maneuver²³ because of its ease of interpretation, its intuitive expansion to a simple sniff, and the difficulty associated with obtaining data from the standard against which tested for the forced expiratory maneuver (discussed in the section that follows).

I propose that I2VE is the first non-invasive, continuous, and dynamically varying technique able to accurately gauge volume status and more importantly fluid responsiveness in both spontaneously breathing and mechanically ventilated patients. To evaluate this hypothesis a series of experiments were designed to compare its performance against a standard of care (Section 3.3.1). Statistical analysis including receiver-operator characteristic curves was performed to survey the results (Sections 3.3.2-3.3.4). Finally a series of conclusions about I2VE's success and limitations were arrived at (Section 3.3.5 and 3.3.7).

²³I am of course speaking only in terms of the spontaneous breathing patient populations here. The specifics of mechanically ventilated patient populations are discussed in Section 3.3.4.




Figure 3.12: An upper limb bioimpedance waveform during forced expiration. Two breaths outward (each limited to 10 cm-H₂O [1 kPa]) have a negative and magnified peak-to-peak amplitude with respect to the normal breathing baseline preceding them.

The work that follows is the first of its kind, though with any luck I hope it is not the last word on the matter.

3.3.1 Experimental evaluation

To test any new technique one must compare it to a standard. In section *Monitoring* of Chapter 1, I reviewed many of the techniques available to assess volume status and responsiveness and a short summary of the facets of these techniques can be found in Table 3.2. Many were found to be poorly predictive (such as central venous pressure), difficult to implement (such as pulmonary artery occlusion pressure), or not currently verified in many types of populations (such as pulse pressure variation). Further still, because the goal of this study was to validate this technique across a broad range of individuals willing and able to give consent within clinical settings, invasive methods were ruled out entirely, leaving only non-invasive techniques at our disposal. Of those, bioimpedance based benchmarks were ruled out due to the possibility of interference from a separate system also injecting current into the body. In so paring we are left with only a handful of ultrasound based techniques to compare against and of those inferior vena cava collapsibility has the most evidence in favor of its use to assess volume status and responsiveness. Respiratory induced variations of the inferior vena cava also have the added benefit of matching many of the facets of the I2VE technique (all except continuous acquisition) and thus allowed for a type of single minimal risk, non-invasive study to be submitted and approved by the University of Michigan Institutional Review Board, streamlining patient consent and data collection.²⁴

²⁴Two separate study protocols were used for this research: *HUM00067675: Noninvasive Monitoring of the Critically Ill and Injured Patient* and *HUM00085550: Assessing Intravascular Volume Status using Limb Bioimpedance Changes*.

Table 3.2: A summary of relevant aspects of current volume status assessment techniques. For a fuller examination of these approaches, see the section titled *Monitoring* in Chapter 1. (CVP - central venous pressure, PAOP - pulmonary artery occlusion pressure, RVEDV - right ventricle end-diastolic volume, LVEDA - left ventricle end-diastolic area, IVC - inferior vena cava, BIVA - bioimpedance vector analysis, SVV - stroke volume variation, PPV - pulse pressure variation, dIVC - change in inferior vena cava diameter, NICOM - non-invasive cardiac output monitoring, typically referring to bioreactance here, I2VE - impedimetric intravascular volume evaluation.)

Method	Mode	Non-invasive	Dynamic	Continuous
Physical exam	-	Yes	No	No
CVP	Pressure	No	No	Yes
PAOP	Pressure	No	No	No
RVEDV	Volume	Yes	No	No
LVEDA	Volume	Yes	No	No
IVC diameter	Volume	Yes	No	No
BIVA	Volume	Yes	No	No
SVV	Volume/Variation	Yes/No	Yes	Yes
PPV	Pressure/Variation	Yes/No	Yes	Yes
dIVC	Volume/Variation	Yes	Yes	No
NICOM	Volume/Variation	Yes	Yes	No
I2VE	Volume/Variation	Yes	Yes	Yes

The approach of the experimental studies that follow was to utilize single-frequency bioimpedance measurements of the upper arm over time. Though multiple frequency spectroscopy has a long history in bioimpedance, for a simple and robust measure of impedance change over time a single frequency approach was chosen (as is often the case in impedance plethysmography). This approach assessing dynamic changes in the peripheral venous volume in response to respiration was then compared to the simultaneous changes in the inferior vena cava diameter. I and my collaborators hypothesized that these dynamic respiratory-induced bioimpedance changes in the upper arm would predict IVC collapsibility and thereby gauge volume status.

Three distinct populations were chosen: (1) normal volunteers and not severely ill or injured emergency department patients; (2) patients undergoing hemodialysis; and (3) those receiving mechanical ventilation in an intensive care unit setting. The first population (normal volunteers and not severely ill or injured patients) was chosen as a proof-of-concept stage of investigation. If the proposed technique was not valid for these simpler cases, it would not bode well for more severe or complicated cases. The second population (patients undergoing hemodialysis) were chosen because of the large volume changes that occur over the course of treatment. It was not uncommon for the subjects of this experiment, many of

whom suffered from sort of kidney failure, to have two, three, or even four liters of excess fluid removed from their system over the course of a few hours. Tracking these rapid intravascular volume shifts with I2VE would prove its validity for a negative fluid challenge case and presumably would do so for the positive fluid challenge. The third and final population (patients receiving mechanical ventilation) represented a case where the intrathoracic pressures driving the signal would be the reverse of the previous two populations. Of the three groups, this was the most difficult to thoroughly test given their vulnerable state and the limits to which changes could be made to the prescribed ventilation routine. Testing across this broad range of patient populations ought to show where this approach excels and where it is limited.

3.3.1.1 Patient recruitment and consent

As previously mentioned the protocols used in this research were approved and monitored by Institutional Review Board of the University of Michigan to uphold medical and research ethics for the benefit, safety, privacy, and well-being of all patients. Before recruitment began within a clinic, physicians and nurses of that clinic were given the opportunity to be briefed on the study's details, the equipment in use, and the potential risks and benefits to their patients. Recruiting was conducted with consultation from the nursing staff of the specific clinic. Non-control subjects for these studies were recruited from the pool of adult patients (18 years of age or older) already present or hospitalized in the University of Michigan Health System. Potential subjects were identified by the research team through either the medical record, their time in the emergency medicine department, their scheduled appointment in the dialysis clinic, or their presence in an intensive care unit. Once identified, subjects were approached in person to obtain consent either from them directly or through a Legally Authorized Representative. To minimize the effects of possible coercion, emphasis was placed on the ability of the subject to disenroll from the study at any point for any reason. Subjects or their Legally Authorized Representatives were given a copy of the research protocol to aid their decisions. Written consent was obtained by a team member prior to any and all study procedures.

Of special concern were the respiratory maneuvers in mechanically ventilated patients where the alteration of respiration carries with it a larger risk than in spontaneously breathing subjects. In this case all the respiratory maneuvers were performed and approved by the respiratory therapy team on a patient-by-patient basis to ensure their health and safety. Furthermore, if at any point any primary caregiver thought our research protocol might negatively affect their patients' outcome, we ceased testing. Though our studies were both non-invasive and carrying minimal risk, the assessment of the caregiver was held in the

highest regard.

The list that follows outlines limitations of use, exclusion criteria, and the separation of data procured by these studies and the clinical decisions made by the treating team. In all instances both medical and scientific ethics were held in the highest regard so as to meet the needs to the individual subject examined and for all subjects for whom this data may one day be relevant.

- Impedance was not measured in limbs arm with either an arteriovenous fistula or dialysis catheter.
- Patients with pacemakers or other similar embedded electronic devices were excluded from examination.
- Vulnerable patient populations including prisoners, children, and pregnant women were excluded from this study.
- A patient may be excluded from examination due to sites of injury or disease preventing measurement (for instance, upper or lower extremity trauma or abdominal injury preventing visualization of the inferior vena cava with ultrasound).
- Study data was not made available to any members of the treating team.
- No clinical decisions made based on either the ultrasound or bioimpedance measurements except in one extreme case where the life of the patient was thought to be in danger by the physician.
- Incidental discoveries were reported to the subjects with the advice that they seek assistance from their primary caregiver.

3.3.1.2 General procedure

Spontaneously breathing subjects were consented, enrolled, and asked to perform a series of respiratory challenges to the best of their abilities. If a subject experienced any pain or discomfort during the task that prevented them from continuing, we immediately ended the examination. Subjects were placed in a supine position and given a series of respiratory challenges that included at least one of the following maneuvers: (1) normal, spontaneous, and/or unlabored breathing; (2) an end-inspiratory and/or end-expiratory hold; (3) a forced inspiration; and (4) a forced expiration. For the forced inspiratory maneuver patients used a threshold inspiratory muscle trainer device (Philips Respironics, Murrysville, PA,) a small respiratory adapter that incorporates flow-independent, one-way valve to ensure a constant

resistance (in this case, between 10-40 cm-H₂O [1-4 kPa] in 2 cm-H₂O increments). A similar device (a threshold positive expiratory pressure adapter also by Philips Respironics) was used to create constant expiratory pressures (between 5-20 cm-H₂O [0.5-2 kPa] in 1 cm-H₂O increments). These devices feature adjustable specific pressure settings and is regularly used in respiratory studies to provide consistent and specific inspiratory and expiratory pressures regardless of how quickly or slowly patients breathe. Subjects were instructed to perform the maneuver as many times as they felt they could for a single exercise. Typically two to four of the forced maneuvers could be done in succession before the subject returned to spontaneous breathing. As discussed in previously, when this was done an average of the sequence was found and used.

Bioimpedance of the limbs was monitored continuously using a tetrapolar configuration via a Biopac MP150 data acquisition system (Biopac Inc., Goleta, CA) and at least one EBI100C electrobioimpedance amplifier (also Biopac Inc., Goleta, CA). Of the four electrodes placed on a limb, two current injecting electrodes were placed far apart and between them two potential sensing electrodes were placed fairly close together (with their centers <5 cm apart). See Figure 3.13 for an example of this on the arm. This arrangement was chosen so that the current density induced by the outer electrodes would be uniform across the region over which the potential was measured.²⁵ A single injection frequency of 50 kHz was chosen as it is a standard value used for these types of impedance plethysmographic experiments (as previously noted). Note further that because the impedance change we hoped to measured was hypothesized to result mostly from blood volume and that blood experiences no lower frequency (α) dispersions (from Table 3.1) that no unintended effects out to have arisen. The current injected was limited by the EBI100C amplifier to a nominal value of 400 μ A RMS of a constant sinusoidal signal. Potential was sampled at a rate of 200 samples/second at the two inner electrodes, sent to the data acquisition unit, which then recorded the value to a computer via AcqKnowledge 4 (Biopac Inc., Goleta, CA).

A standard curvilinear or phased array ultrasound probe was used concurrently to visualize the inferior vena cava in a subxiphoid view via a commercially available device (Mindray M7, Mindray North America, Mahwah, NJ). The inferior vena was examined with patients in the supine position subcostally in a longitudinal view by one of four emergency medicine physicians with special training and credentialing in sonography. Video clips of the ultrasound (commonly referred to as cineloops) over the length the respiratory maneuvers (10-15 seconds in length) were recorded and saved in the DICOM format to be processed via the algorithm presented in Chapter 2.

Bioimpedance monitoring of the limb and ultrasound imaging of the inferior vena cava

²⁵The results of equations 3.33 and 3.46 make this desire pretty clear.




Figure 3.13: Electrode placement along the upper limb for the I2VE experiments reported here. Two current injecting electrodes reside near the shoulder and the wrist/forearm and two potential sensing electrodes were placed closer to each other along the bicep. Image from US 2015/0031966 A1 and is used for educational purposes.

were performed simultaneously for all respiratory maneuvers.

3.3.1.3 Special considerations for mechanically ventilated patients

Because patients receiving mechanical ventilation are in a particularly sensitive position, special care was taken to ensure their health and safety. The design of these specific respiratory maneuvers was done in consultation with respiratory therapists. The specific maneuvers included (1) non-perturbed mechanical ventilation (no ventilator changes), (2) an inspiratory hold up to 10 seconds if the subject could tolerate it, (3) an expiratory hold up to 15 seconds if the subject could tolerate it, (4) a mechanical sigh of 1.5x the tidal volume, (5) a transient increase of tidal volume to 10 cc/kg for up to 2 minutes if the subject could tolerate it, and (6) a passive leg raise. Each of these procedures causes a reversible change in blood volume present in the thorax, the extent and duration of which is procedure dependent.

None of these maneuvers has been associated with any deleterious effects. In fact, each has been used clinically to different ends in the assessment of respiratory and hemodynamic parameters. Inspiratory holds and sighs are commonly used as strategies to recruit alveoli[208, 209] as a result of the widespread use of low volume lung protective strategies (7-8 cc/kg of tidal volume).[210] Expiratory holds have been used to assess volume

responsiveness in mechanically ventilated patients with responsive subjects experienced increased cardiac output and those non-responsive patients experiencing no effect.[211, 212] Transient increases in tidal volume (to 10 cc/kg) have been used to assess volume status via inferior vena cava distensibility [68, 70]. As previously mentioned the passive leg raise maneuver has been used to great effect to produce a reversible, noninvasive fluid challenge to the central cardiovascular system by increasing venous return. The specifics of this procedure are outlined in Section 1.6.6.4. All respiratory maneuvers were performed at two to three minute intervals to allow the patient to return to baseline between challenges. A physician faculty investigator was present over the course of all trials and all maneuvers were approved and performed by respiratory therapy team.

Early termination would result if a Legally Authorized Representative revoked permission, if a member of the care team needed to do a procedure in which the presence of the study team or its equipment would interfere, if the respiratory therapists deemed that the patient was intolerant to all respiratory challenges, or if a patient’s vital signs (heart rate, blood pressure, oxygen saturation, etc.) varied by more than 10% over the course of the trial.

3.3.1.4 Data and statistical analysis

Raw bioimpedance data for each respiratory maneuver was initially subjected to an iterative simple moving average²⁶ (three passes through a 20 point moving average) to smooth. To separate the respiratory and cardiac signals during spontaneous breathing, the data was further low-pass filtered (with a corner frequency of 0.6 Hz) to find the respiratory signal. This portion of the signal was then removed from the all-pass signal to yield the cardiac signal by the principle of superposition. Over the course of the respiratory challenge, the peak-to-peak amplitude of the impedance signal was identified and is reported throughout as Z_{pp} and is equivalent to the ΔZ of equation 3.28 in this context. Because the individual contributions of the cardiac and respiratory impedance signals are difficult to discern in practice (forced maneuvers causing frequency contributions at or above the level of the heart rate) and in theory (separating the arterial and venous components is an as of yet unsolved problem), Z_{pp} refers to the peak-to-peak amplitude of the entire impedance waveform over the course of a single breath and not exclusively to either the cardiac or respiratory component

²⁶Defined here and elsewhere as applying to a sequence consisting of M points of x ($\{x_i\}_{i=1}^M$) to be

$$s_i = \frac{1}{n} \sum_{k=1}^{i+n-1} x_k \quad (3.47)$$

(though the separation and analysis of the two is discussed later).

In this study, Z_{pp} was of two forms and it was their ratio that was the basis of the I2VE parameter, dZ , representing and scaling the difference in impedance caused by forced respiration. The first was a baseline value, Z_b , which was chosen to be normal, unguided breathing in unventilated subjects and the unchanged therapist prescribed settings in ventilated subjects. The second Z_{pp} was one of the forced respiratory maneuvers, Z_m , previously discussed. Since over the course of a single exercise a subject may have performed the maneuver any number of times, the average Z_m value was used in the calculation of dZ via

$$dZ = \frac{1}{Z_b} \sum_{i=1}^n \frac{Z_{m,i}}{n} \quad (3.48)$$

where i is the counter variable and n is the total number of maneuvers in an exercise. Though not used in this analysis, sometimes a change in the magnitude and offset of the baseline Z_{pp} was observed, defined here as $Z_{b,a}$ and $Z_{b,s}$, respectively. Figure 3.14 shows a typical example of an exercise with each of these variables listed.

All inferior vena cava measurements in B-mode were made after identifying the point of maximal collapse, approximately 2 cm caudal to the hepatic vein inlet, using frame-by-frame analysis to identify the point of maximal and minimal diameter using the algorithm of Chapter 2. Measuring at this location and in this way has been previously validated both in our own experience (see Section 2.4.2) and in the literature [159]. The maximum and minimum diameters were used to calculate the collapsibility index as described in equation 2.110. Although some sonologists have studied M-mode imaging to calculate inferior vena cava collapsibility and distensibility [154] others have suggested (and I have seen through personal experience) that respiration, especially forced respiration, results can result in significant displacement of the IVC which may make M-mode measurements less accurate [159, 160]. In those cases where the automatic assessment of the ultrasound image via the methods described in Chapter 2 could not be accomplished, manual examination by a specially trained and credentialed ultrasound physician was performed.

Ultrasound and bioimpedance data analysis was performed without prior or simultaneous knowledge of their results. Once all necessary measurements were made the two could be compared. To test the direct relationship of dZ to the caval index, $dIVC$, a nonlinear least squares regression was performed (via a Trust-Region algorithm) on the entire set of data to a function of the form

$$dIVC = A \ln \left(\frac{Z_m}{Z_b} \right) + B \quad (3.49)$$




Figure 3.14: A typical impedance waveform observed before and after a forced respiratory maneuver demonstrating each of the the I2VE parameters: the baseline peak to peak impedance change due to spontaneous breathing, Z_b ; a series of forced inspiration causing individual amplitude changes, $Z_{m,1-3}$; a potential change in the baseline amplitude after the maneuver, $Z_{b,a}$; and a possible baseline shift, Z_{bs} , increasing the impedance minimum, distinct from the amplitude change. The I2VE parameter, dZ , may be calculated from equation 3.48. Note that Z_b here is shown as a single wave for convenience. In reality the unperturbed breathing baseline is an average value of many waveforms.

where A is the slope of the line and B is the offset at $dZ = 0$.

Receiver operator characteristic (ROC) curves for single and multiple thresholds were used to evaluate the diagnostic performance of the technique. That is, to assess the accuracy of I2VE to discriminate fluid responsiveness from non-responsiveness, the true positive rate, known as sensitivity, was plotted as a function of the false positive rate, represented as 1-specificity. In general, the area under this curve is a measure of how well the diagnostic technique is and in our case represents how well I2VE predicts volume status. An area under the curve (AUC) equal to 1, indicates a perfect ability to distinguish true positives from false positives, such that an incorrect positive prediction never happens.²⁷ For single threshold detection, the value of dZ was used to predict whether a $dIVC$ would lie either above or below a threshold of collapse (for instance, to evaluate how well dZ is able to

²⁷For a system with two choices (say, heads or tails), random guessing would yield an AUC equal to 0.5 by producing a diagonal line from the origin to 1 in the ROC graph. This line, known as the line of non-discrimination, is about as poor as a predictor can get and why coin flips for medical diagnosis should be avoided.

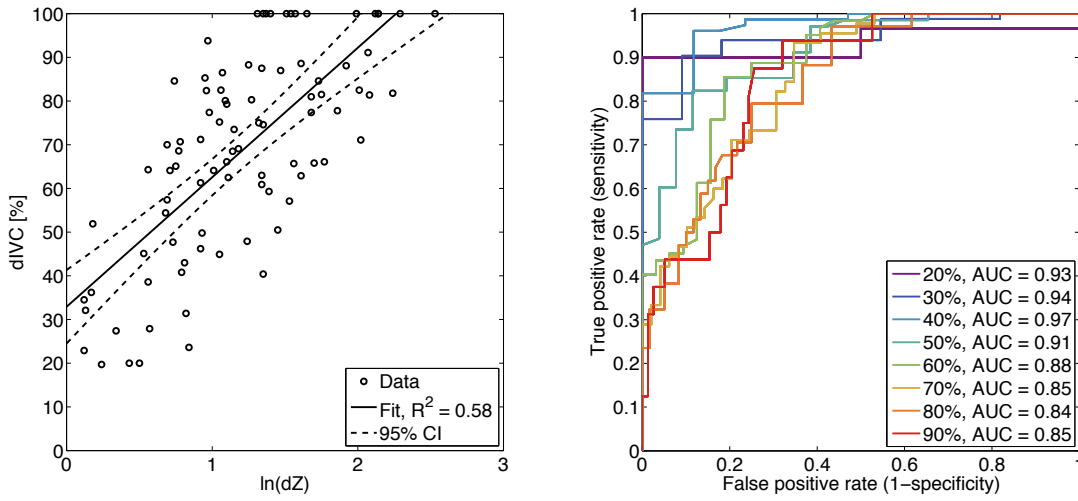
predict $dIVC > 40\%$, relevant for fluid responsiveness [72].) For the multiple threshold detection, dZ was used to predict whether a value was below a lower bound or above an upper bound (for instance, $dIVC < 30\%$ vs. $dIVC > 50\%$). When the lower and upper bounds are the same, the multiple threshold analysis is the same as the single threshold analysis. In all cases, the AUC was treated as the measure of diagnostic accuracy.

Since this study was not a traditional outcome testing study, it is not possible to determine a discrete power. However, a population of greater than 30 subjects with complete data in each case was deemed sufficient by the team (including a member who is a biostatistician) to test the hypothesis that bioimpedance is a suitable surrogate for inferior vena cava collapsibility. Furthermore, the three populations were targeted to provide sufficient exposure to a range of dynamic changes in physiology and therefore enough signal variance to utilize the above stated techniques.

3.3.2 I2VE of control and hospitalized subjects


The first test of feasibility for this technique was performed on spontaneously breathing subjects of two types: normal “healthy” volunteers recruited from the general population and non-critically ill patients present in the emergency room. Given the focus of this first study, the correlation of limb impedance to inferior vena cava diameter, the two populations were not considered separately as there was no reason to believe the cardiovascular or pulmonary functioning of the two were distinct. And indeed the team’s observations bore this out. A total of 40 subjects were enrolled for this study. However, the inferior vena could not be imaged on two patients and their results were excluded. Of the remain 38 subjects, the demographic breakdown was as follows: 20 were male, 18 female; 25 were white, 8 were black, and 6 were Asian; the average age was 41.2 years with a standard deviation of 20.2 years; and the average weight of the subjects was 75.0 kg with a 18.2 standard deviation.

The results of this population can be seen in Figure 3.15. Figure 3.15(a) plots the caval index, $dIVC$, against the change in impedance caused by the forced maneuver normalized to the baseline, dZ . A mild positive linear relation was observed ($R^2 = 0.58$) for the linear regression model $dIVC = 29.7 \ln(dZ) + 32.9$. Figure 3.15(b) shows single threshold ROCs for $dIVC$ collapses from 20-100% with AUCs in these instances varying from 0.84 to 0.97. The AUC of each collapse from 20-100% in 1% increments can be seen in Figure 3.15(c). A peak AUC (of 0.971) with a single threshold can be seen at 40%. The AUC of the ROC curves found via multiple threshold analysis is plotting as a function of lower and upper collapse thresholds in Figure 3.15(d).




(a)

(b)



(c)



(d)

Figure 3.15: The results of the volunteers and emergency room patients: (a) inferior vena cava diameter collapse vs. impedance change from baseline; (b) ROC curves of several single threshold tests; (c) area under the ROC curves for every collapse threshold observed (20-100%); and (d) area under the ROC curve as a function of the lower and upper bounds of the multiple threshold analysis.

Though there is a correlation between dZ and $dIVC$, it was not as great as one might have hoped. That is, one can not simply declare that a certain impedance value, even if normalized to a baseline, corresponds directly to an inferior vena cava diameter. However, when determining whether $dIVC$ is above or below some threshold, great predictive power was observed, even at a single threshold. The I2VE technique yielded an area under the ROC curve equal to 0.971 at 40%. This result matches well with that reported by Muller et al. [72] as the index of high sensitivity and with Barbier et al. [154] as the average value of inferior vena cava collapse for volume responsive subjects (true of this population). One of the most striking results is the technique's ability to distinguish results below a lower threshold and above an upper threshold. In many cases AUCs of 1 were observed. Of particular diagnostic importance is the resolving power seen with lower bounds between 20-40% and upper bounds of >40%, as this a region known to indicate volume response and responsiveness.

3.3.3 I2VE during hemodialysis

Believing that the results of the previous section boded well, it was then necessary to test I2VE's ability to monitor rapid intravascular volume shifts. To do so, a population where large, controlled volume shifts were performed regularly was chosen, namely subjects receiving hemodialysis. A total of 52 patients undergoing hemodialysis as a course of treatment were consented, enrolled, and evaluated. Of those, the inferior vena cava could not be visualized in five people, and thus only 47 were evaluated to completion for the results presented here. The demographics of this population were as follows: 28 were male, 19 female; 24 were white, 21 were black, and 2 were Asian; the average age was 58.1 years with a standard deviation of 12.2 years; and the average fluid removed as 1.9 L with a standard deviation of 0.9 L.

Subjects were evaluated "before" and "after" dialysis, here defined to be within half an hour of the start and end of treatment, respectively. Figure 3.16 shows the relationship between the inferior vena cava collapse and the natural log of the impedance changed normalized to a spontaneously breathing baseline both (a) before and (b) after dialysis. A linear regression model was fit for both distributions producing a model of $dIVC = 35.9 \ln(dZ) + 16.2$ for the pre-dialysis population and a model of $dIVC = 41.1 \ln(dZ) + 13.8$ for the post-dialysis population. No statistically significant difference in these models was observed ($p = 0.15$) as a function of time.

As a check on using inferior vena cava collapsibility as a standard against which to evaluate the large volume shifts of hemodialysis, I analyzed its dimensions for each subject

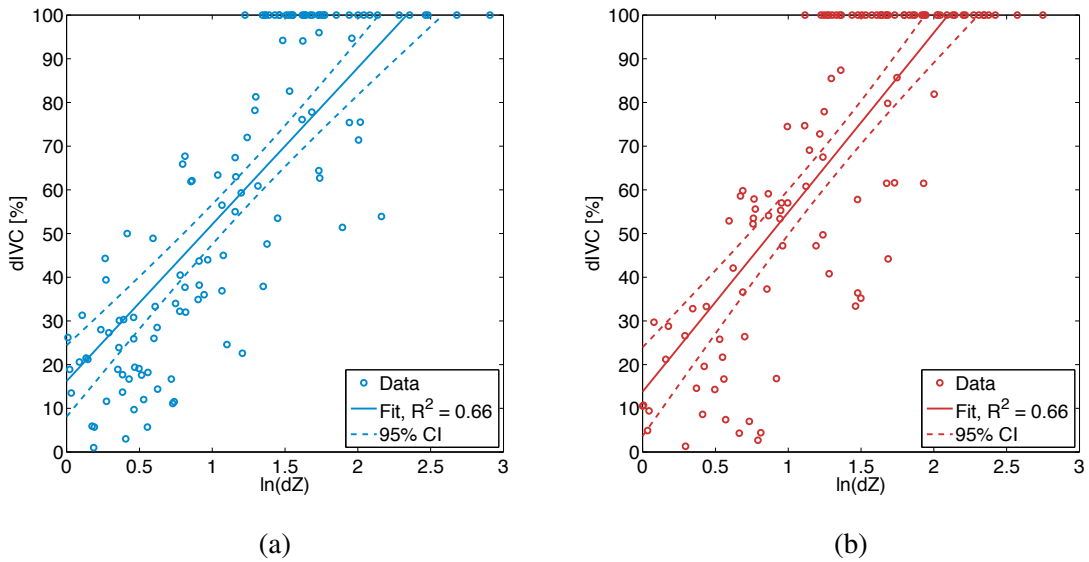


Figure 3.16: Inferior vena cava diameter collapse vs. impedance change from baseline (a) before and (b) after dialysis. Both the before and after measurements show moderate correlation (R^2 values of 0.664 and 0.657) for their linear regression models of $dIVC = 35.9 \ln(dZ) + 16.2$ and $dIVC = 41.1 \ln(dZ) + 13.8$, respectively.

under normal breathing conditions before and after dialysis and found they changed significantly (as determined by a paired, one-tailed t-test in all cases): the maximum diameter went from 17.7(7.4) mm to 15.2(7.7) mm ($p = 0.0002$); the minimum diameter decreased from 13.0(9.1) mm to 9.75(8.9) mm ($p = 0.0001$); and $dIVC$ increased from 41.2(32.3)% to 52.4(34.0)% ($p = 0.002$).²⁸

ROC curves evaluating the sensitivity and specificity of the I2VE parameter, dZ , for a subset of single thresholds of inferior vena cava collapse (10%, 30%, 50%, 70%, 90%) before and after dialysis are shown in Figure 3.17(a) and (b). In both cases, the extrema have the lowest resolving power (10% and 90% both before and after have AUCs of 0.88 and 0.93), while the peak accuracies lie somewhere in between (at 50% before (AUC = 0.98) and at 70% after (AUC = 0.96) in this figure).

Figure 3.18(a) and (b) shows the AUC of each ROC curve for collapse thresholds from 10-100% before and after dialysis. Here, as in Figure 3.17(a), a peak at 0.98 exists at the 50% collapse threshold in subjects before dialysis. After, two potential peaks of 0.975 and 0.959 appear at 33% and 61%, two possible regions of high accuracy in this population. In both situations, the AUC of the ROC curve is ≥ 0.90 for collapse thresholds greater than

²⁸Though a few researchers have performed similar studies [213, 214] in hemodialysis settings, this technique is not a current standard, and thus this check was useful to validate our intuition.

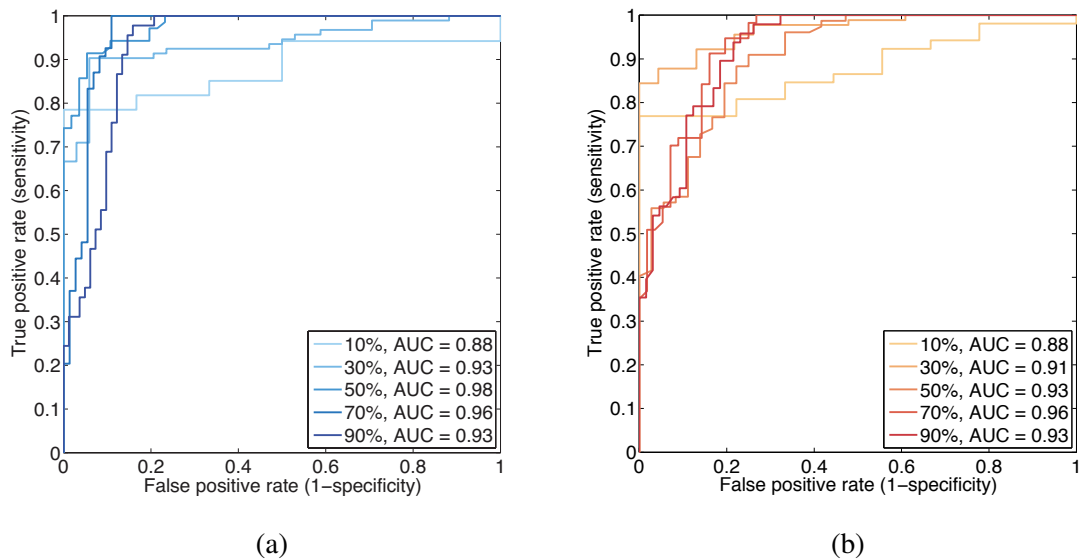


Figure 3.17: Example ROC curves of five thresholds of inferior vena cava collapse (10%, 30%, 50%, 70%, 90%) (a) before and (b) after dialysis. Areas under the curve in these examples ranged from 0.88 to 0.98 before dialysis and from 0.88 to 0.96 after, demonstrating good predictive power in both scenarios.

20%.

Using dZ to distinguish between lower and upper thresholds of inferior vena cava collapse before and after dialysis produces the results seen in Figure 3.19. Pre-dialysis, a large portion of the collapsibility space marked off by the region encompassing lower thresholds from 10-60% and upper thresholds from 30-100% has AUCs of ≥ 0.95 , with a significant subset equal to 1.00 (for example, 20% vs. 50% and 40% vs. 70%). An interesting phenomenon can be seen in the post-dialysis population where two distinct regions—one bounded by a 10-36% lower threshold and 34-100% upper threshold, the other bound a 10-60% lower threshold and 60-100% upper threshold—each exhibit higher accuracy (as measured by AUC) than those thresholds lying outside those bounds. Indeed, those collapse thresholds within the overlapping regions from the 10-36% lower limit and 60-100% upper limit have AUCs equal to 1.00.

Figure 3.20 shows the combined results of all subjects pre- and post-dialysis. The linear regression model fit to the total data (seen in Figure 3.20(a)) was found to be $dIVC = 38.08 \ln(dZ) + 15.48$ with a moderate correlation coefficient, $R^2 = 0.69$, slightly higher than either the pre- or post-dialysis results. Figure 3.20(b) shows ROC curves testing the accuracy of the I2VE parameter, dZ , for a subset of single inferior vena cava collapse thresholds (10-90%). The trend here is similar to that observed in Figures 3.18(a) and (b)

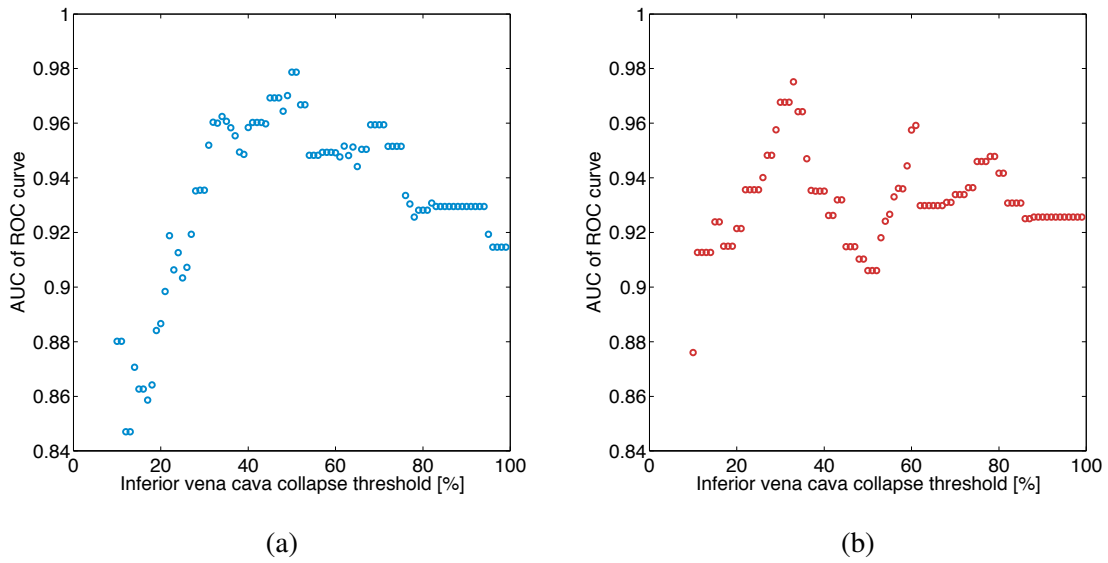


Figure 3.18: Area under the ROC curves for collapse thresholds from 10-100% in subjects both (a) before and (b) after dialysis. A peak of 0.979 was observed at 50% collapse before dialysis, while two potential peaks of 0.975 and 0.959 appear at 33% and 61% after.

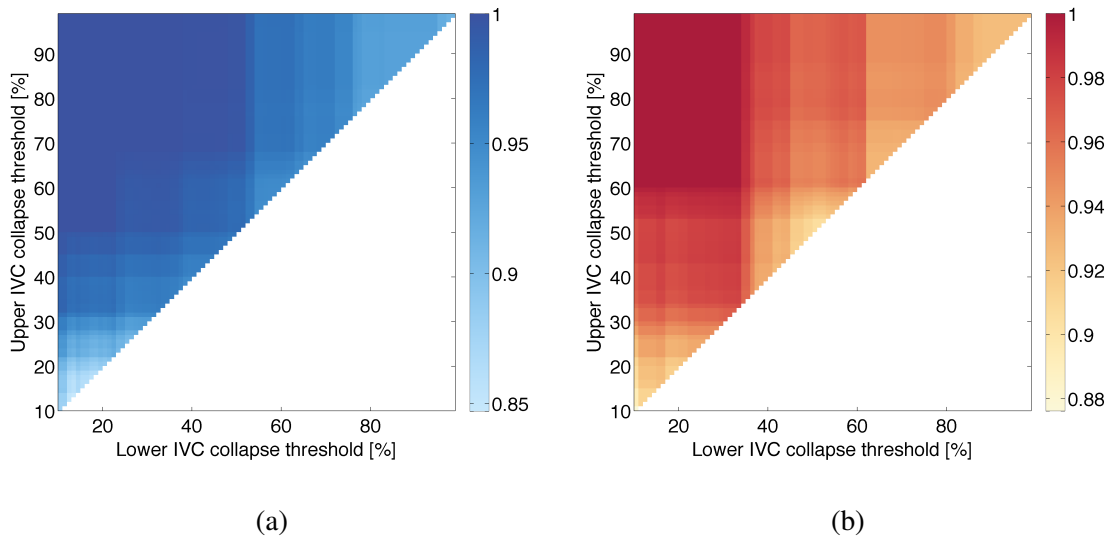


Figure 3.19: Areas under the ROC curve as a function of the lower and upper collapse threshold (a) before and (b) after dialysis. Large sections of both graphs show AUCs greater than 0.99, with the extrema seeing the lowest values (still above 0.85 in all cases before dialysis and above 0.88 in all cases after).

where extreme values of collapse produce ROC curves with the lowest AUCs (in this case, for a low of 0.865 at 10%), while the peak (of 0.968 at 30%) resided in between. In point of fact, as can be seen in Figure 3.20(c), the peak AUC of a single threshold occurs at 33% with a value 0.954. The multiple threshold analysis of Figure 3.20(d) shows a region a high AUCs (≥ 0.95) extending from the lower boundary of 0-33% collapse to the higher boundary of 33-100% collapse.

The total results encompassing both the pre- and post-dialysis scenarios suggest I2VE is particularly suited to distinguish between inferior vena cava collapses above or below 33% and to characterize the extent of collapse on either side of this threshold with remarkable fidelity. That this value so aligns with the average found in the spontaneous breathing population of Section ?? should not go unnoticed. In fact, resolving down to the tenths of a percent for collapse indicates the true peak AUC occurs at 33.4% with a value of 0.972, which is only 0.2% lower than that reported by the algorithm on the 47 subjects of Section ?. Though this agreement at this level of precision may be coincidental, the physiological relevance is likely not.

3.3.4 I2VE in mechanically ventilated subjects

Mechanically ventilated subjects were chosen to test the efficacy of I2VE in intensive care setting. A total of 42 patients receiving positive pressure ventilation as a course of treatment were consented to be a part of this study. Because of the sensitive nature of the population to respiratory exercises and some difficulties obtaining views of the inferior vena cava, only 31 subjects had complete data sets to review and it is those that are presented here. Caution should be taken when interpreting these results as there is an inherited selection bias from this paring. Mindful of these limitations, I believe this data is valid enough to report here.

Adding to the difficulties of validating I2VE within this population, positive pressure ventilation inverses the relationship between inferior vena cava diameter and respiration. Instead of collapsing during inspiration, the inferior vena cava distends, since a breath is pushed into the lungs, instead of pulled by a patient. Therefore a more appropriate measure of the caval index in this case is to normalize the difference in diameters with respect to the minimum (rather than the maximum, as was the case in the spontaneous breathing population trials). To this end, I define the positive difference in the inferior vena cava, *pIVC* as

$$pIVC = \frac{D_{max} - D_{min}}{D_{min}} \quad (3.50)$$

where D_{max} and D_{min} are the maximum and minimum diameters of the inferior vena cava, respectively. Thus defined, Figure 3.21(a) shows the relationship between *pIVC* and *dZ*.

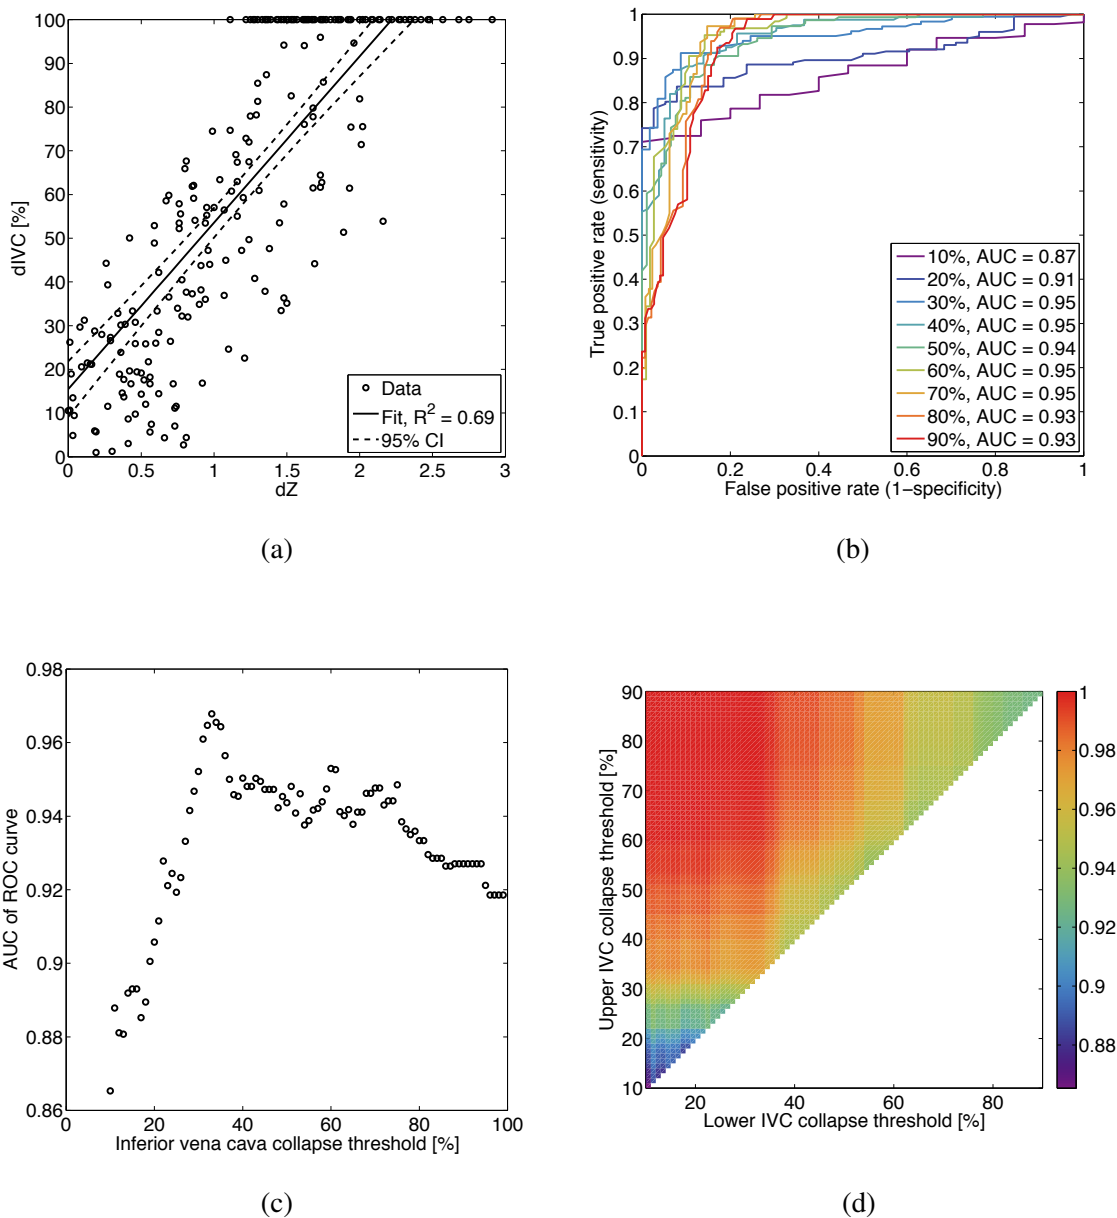


Figure 3.20: The total results from patients receiving hemodialysis (both before and after treatment): (a) inferior vena cava diameter collapse vs. impedance change from baseline (showing moderate correlation ($R^2 = 0.69$) for its linear regression model of $dIVC = 38.08 \ln(dZ) + 15.48$); (b) ROC curves of several single threshold tests (with AUCs ranging from 0.87 to 0.95); (c) area under the ROC curves for collapse thresholds from 10-100% (with a peak value of 0.968 at 33% collapse); and (d) area under the ROC curve as a function of the lower and upper bounds of the multiple threshold analysis (with the region between the lower thresholds of 10-30% and upper thresholds of 30-100% yielding the best results (AUCs > 0.95)).

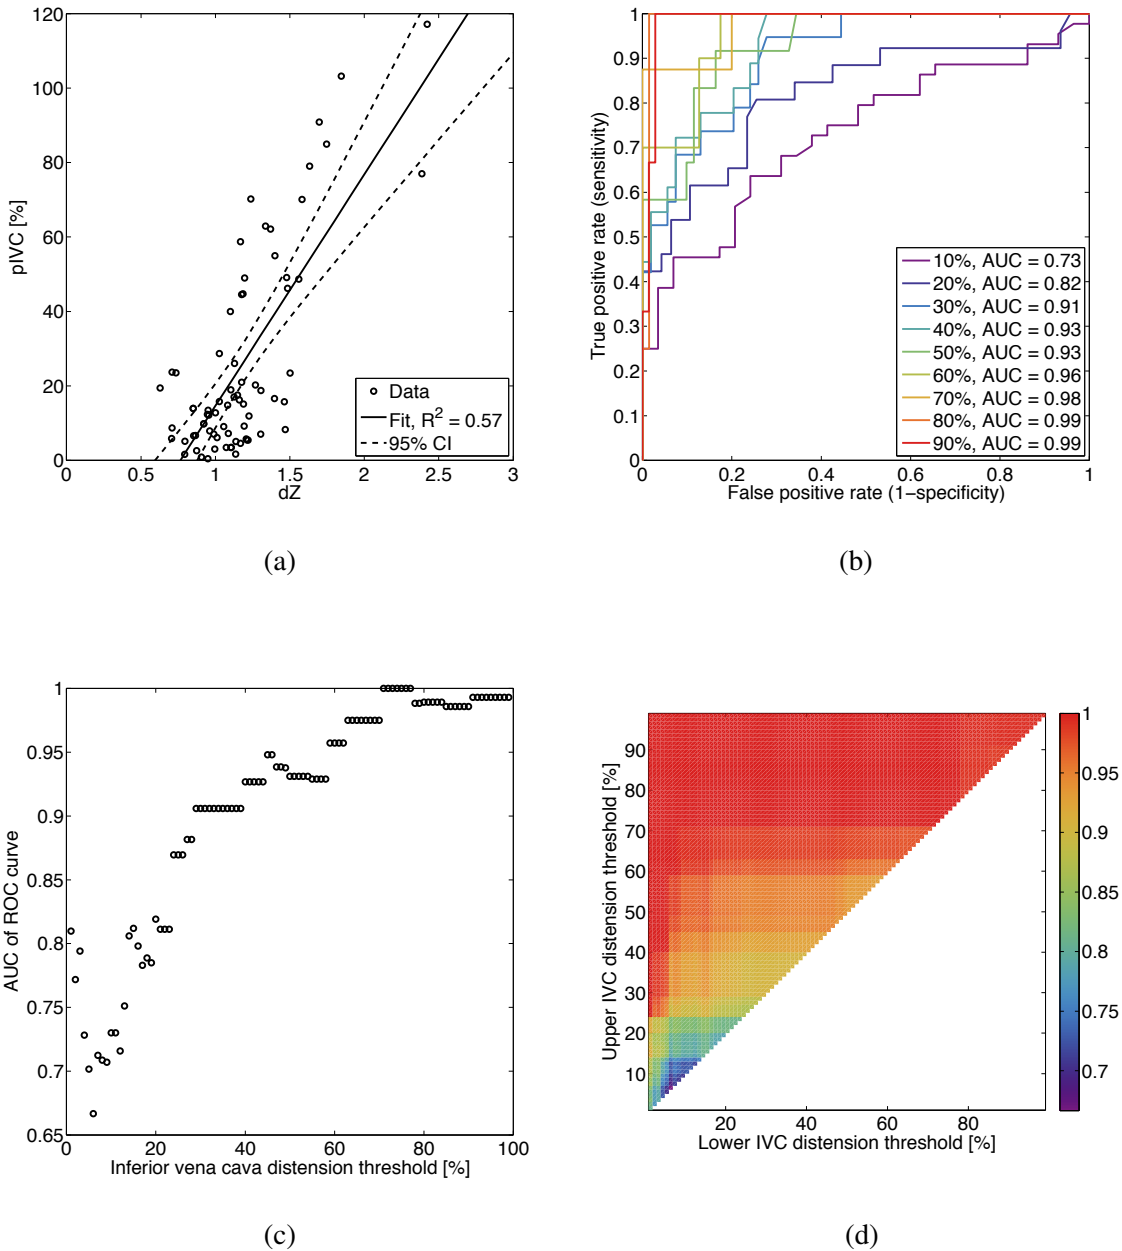


Figure 3.21: The total results from patients under mechanical ventilation: (a) inferior vena cava diameter distension vs. impedance change from baseline (showing moderate correlation ($R^2 = 0.57$) for its linear regression model of $dIVC = 62.15(dZ) - 47.49$); (b) ROC curves of several single threshold tests (with AUCs ranging from 0.73 to 0.99); (c) area under the ROC curves for distension thresholds from 1-100% (with a dip to 0.67 at the 6% threshold); and (d) area under the ROC curve as a function of the lower and upper distension bounds with AUC increasing as a function of increasing of the upper limit threshold.

With a median inferior vena cava distension of less than 15%, very little variation was observed and is quite clear in this figure, where the lower end of the range outweighs the higher end in sample density. This result carries over to both Figures 3.21(b) and (c) where the resolving power of the technique is weakest at the lower range of *pIVC*. Likewise, this trend carries over into the multiple threshold analysis where the area under the ROC curve increases predominantly as a function of the upper limit. Whether this is an actual predictive function or one merely caused by an imbalanced sample must yet be determined.

3.3.5 A few limitations of I2VE

Though I believe I have demonstrated impedimetric intravascular volume evaluation's unprecedented capabilities—non-invasive, dynamic, and continuous volume status monitoring with a higher correlation to inferior vena cava collapsibility than any other technique reported in the literature—as currently reported, a few limitations exist.

The first stems from the use of inferior vena cava collapsibility itself as a standard against which to compare. Though I feel justified in having used this approach (and those justifications can be seen in Section 3.3.1) given the other options, I have my reservations on its direct applicability to I2VE. One of my key concerns, as of yet unaddressed in the field, is the normalization of collapse to a maximum diameter without accounting for what that diameter on its own might indicate. As an example an inferior vena cava diameter going from 20 mm to 10 mm produces the same collapsibility as one going from 10 mm to 5 mm, though if occurring within the same person represent two different physiological states. This was not an uncommon occurrence, especially in hemodialysis subjects where the venous system may have been abnormally distended from volume overload. A relative scaling coefficient to normalize the range or using two parameters to characterize the diameter and its change over time could possibly be used to correct for this.

A similar critique can be made of my normalization of the forced maneuver to a normal breathing baseline. Patients with large peak to peak amplitudes during normal respiration as a result of pre-existing volume depletion are unlikely to evince significantly larger amplitudes in response to a forced maneuver. The converse is always true as in the case massively overloaded patients whose peak to peak impedance amplitude during normal and forced respiration is exceedingly small. That is, when volume status is already at an extreme minimum (and inferior vena cava collapse is at 100% through much of the respiratory cycle) or an extreme maximum (and the inferior vena cava barely wavers from its maximum diameter), respiratory perturbation might not manifest in a measurable way (thus violating the third assumption of Section 3.2.2). One way to correct this may be

to implement one of my suggestions for the inferior vena cava diameter characterization, though to do so, one must decide how they will treat the respiratory and cardiac components of the signal—separately or as part of a total signal. The reasons for doing so are twofold. First, the cardiac signal comes from the arterial side of the cardiovascular system and the respiratory signal comes from the venous side and are thus reflect and are affected by volume status in different ways. Second, the proportion of the cardiac component is not constant across subjects and vary wildly in my investigations from 5-90% of the overall signal (though typically residing close to 10-20%). As such future normalization, scaling, or measurement ought to account for the presence and proportion of the cardiac component.²⁹

One final limitation to the technique more broadly is the need to perturb the system to gauge the volume status. Although several traditional respiratory maneuvers could be used (such as the valsalva or a sniff) in many settings, there are situations when the information would be useful but difficult to obtain (such as when a patient is sleeping or during an emergency). Choosing a different baseline against which to scale the dynamic respiratory changes could allay these needs.

3.3.6 Other observations

Even with the addition of I2VE as demonstrated in the previous sections, the full extent of impedance plethysmographic techniques for hemodynamic assessment has been far from wholly realized. Aside from the simple yet fundamental ability to measure vital signs such heart rate and respiratory rate, I have observed many other phenomena during my investigations worth future exploration (as outlined in Section 3.3.7). Two such phenomena that I have seen in enough detail are reported below. The first, the arterial component of the signal, has been utilized by other researchers for decades, though not to the extent I would call for. The second, the depth and rate of respiration as measured via impedance plethysmography, has not to my knowledge been documented elsewhere.

3.3.6.1 Arterial contributions

Let us recall that the impedance plethysmographic waveform represents the changing volumes of conductors within an limb and that over a short period of time blood volume changes are the most significant. For I2VE, an emphasis was placed on the venous contribution because of its role in intravascular volume assessment, which manifest prominently

²⁹Because this technique relies on the *change* in impedance and that to be valid that change must be small relative to the total impedance value, the actual impedance value does not matter. This is convenient since electrode type and placement can affect the total value via an offset, but does not in most cases scale the difference.




Figure 3.22: The arterial and venous components of impedance plethysmography. (a) The original signal and the respiratory signal, the latter of which represents venous blood movement. (b) The cardiac signal of impedance separated from the original and respiratory signals. Note the relative amplitude is approximately $\pm 10\%$ that of the total signal. (c) The first order derivative of the cardiac impedance signal over time, an approximate stroke volume over time measurement.

in the respiratory signal. However, as can be seen in Figure 3.10, a cardiac variation was observed, likely stemming from an arterial contribution.

Figure 3.22(a) shows another typical impedance waveform equivalent whose arterial contribution was separated from the respiratory signal in Figure 3.22(b). In this latter figure was found by subtracting the respiratory signal from the original signal, producing an impedance over time graph of the cardiac signal. Taking the first order derivative, as intimated in equation 3.29, yields a tracing of the change in impedance as a function of time, $\Delta Z/\Delta t$ in Figure 3.22(c), a measure of cardiac output [195, 215, 197].




Figure 3.23: Representative arterial waveform found via impedance plethysmography by taking the first order derivative of the cardiac signal. Characteristic features seen in other arterial monitoring methods, such as systolic uptake, systolic decline, and the dicrotic notch are present here.

Though such an approach has been implemented several times for thoracic impedance, I have seen no study of the peripheral equivalent. I discuss the ramifications of such a measure in Section 3.3.7.

Although the results reported here used non-perturbed breathing exclusively as the baseline against which to compare the respiratory maneuver, some initial study on my part suggests that to me the arterial component of signal could be used in its place.

Isolating an individual heart beat, as shown in Figure 3.23, we can see the $\Delta Z/\Delta t$ waveform's similarity to other arterial measures such as stroke volume and pulse pressure. Such important diagnostic features as systolic uptake, peak systole, systolic decline, the dicrotic notch, and diastolic run off are all present in the order and magnitude one would expect (note the resemblance to the aortic pressure signal in Figure 1.3). As such, one could possibly leverage the large body of arterial waveform analysis to further extend the studies presented here, possibly even stroke volume variation methods (such as those discussed in 1.6.6.1). Moreover, if the constraints of the monitoring environment were adequately accounted for, stroke volume could be assessed accurately in combination with the heart rate, quantifying cardiac output to the limb (see Figure, 1.4).




Figure 3.24: Impedance over time of a sleeping, snoring patient. The subject takes a series of large and rapid breaths before settling down towards a shallow breathing baseline, only to start the cycle anew.

3.3.6.2 Respiratory function

Given that the reliance of I2VE on respiration, it should come as no surprise that respiration itself can be monitored well. Both the frequency and intensity of breaths affect intrathoracic pressure which in turn alters venous return, itself able to be tracked via I2VE. Figure 3.24 shows the impedance plethysmographic tracing of a sleeping individual whose breathing becomes more shallow over time, causing them to take a series of deep inspirations to recover before settling back to shallow breathing. This pattern continued for nearly an hour and a half for this particular individual (this was during their dialysis treatment), periodically settling into a stable breathing rhythm before lapsing into it again.³⁰ Tracking the respiratory rate from the breath to breath extrema, as seen in Figure 3.25, shows the steady decrease in the rate after the initial deep inspiration to the shallower. Thus, this example showcases the unique ability of impedance plethysmography to track the presence, extent, and rate of respiration, making




Figure 3.25: Decreasing respiratory rate as measured from the second half of Figure 3.24. Being able to track the breath to breath accuracy with impedance could open up many potential avenues of research.

³⁰The teams' notes indicate that this patient was snoring heavily during the times this was recorded.

possible the detection of variety of breathing states from eupnea to apnea, from hyperpnea to hypopnea.

3.3.7 Conclusions and the goals of future work

In this chapter I hope to have convinced the reader that impedimetric intravascular volume evaluation, a bioimpedance based technique that treats the peripheral vasculature as an extension of the inferior vena cava, can detect the intravascular volume status and responsiveness of a patient with unprecedented levels of accuracy, unobtrusiveness, and ease. Four electrodes placed on an upper limb with the right combined physiological and signal processing approach have yielded better results than every other technique outlined in Chapter 1, with AUCs for ROC curves well above 0.95 in the vast majority of scenarios. Such a degree of precision for a continuous and non-invasive volume status monitor has not to this point been shown in the literature, though the clinic is in desperate need of just such a method.

We have, however, seen that some edge cases—those of massive overload, volume depletion, or where the cardiac signal may dwarf the respiratory signal—could potentially pose problems for I2VE in the future if unperturbed breathing is alone treated as the baseline against which to scale forced respiration. I have suggested a number of ways to solve this including normalizing the cardiac signal or choosing a multi-parameter measure (and thus a new normalization framework). Whether these are effective remain to be seen. However, the greatest sensitivity and specificities results are around the areas of interest for conditions such as intradialytic hypotension and its ilk, and thus might be sufficient as is for many clinical applications.

The final results I have shown demonstrating the ability to reproduce the arterial waveform point to a future implementation wherein both cardiac output and venous return could be monitored via a single device to find the actual operating point of the heart (Figure 1.5) for the first time. Combined with a respiratory monitor capable of gauging not only the rate but the intensity of respiration, and one would have a device capable of measuring nearly ever vital sign physicians currently look to as a matter of standard practice. An all in one medical monitor has been the dream of many engineers and researchers. In the chapter that follows I hope to show at least a rough outline of what I suspect such a device might look like.

CHAPTER 4

Wearable bioimpedance monitoring

Health lies in action.

Will Durant, *Fallen Leaves*

4.1 The rise of wearable monitoring

Electronics have gotten ever smaller, their capabilities ever greater. The simultaneous rise of wireless communication with the miniaturization of computation has led to a boom throughout many industries to create small, portable, powerful devices. Among those devices are a broad range of sensors from the chemical to the optical, from the mechanical to the electrical. These sensors often enable faster and more accurate measurements than their larger more energy-hungry predecessors. We can measure more signals in more places than ever before. And those of our physiology are no exception.

I believe it is my job as a biomedical engineer to meet the needs of the patient when and where they have them. To that end, the last part of my work was to develop a wearable monitor to measure the intravascular volume status of a person building atop the foundation laid by the previous two chapters.

4.1.1 Biomedical signals in context

Physiology happens beyond the bedside. How our bodies work and fail is often a function of how and where they find themselves in the environment. Sedentary lifestyles, air pollution, the weather, what one ate a week ago, many factors lie within an individual's slice across this world. Traditionally, however, the only picture a caregiver could paint of the patient before them were their medical signals at the bedside and the record of such signals scattered as lone islands in a subject's history. Recent developments have begun to bridge these islands, to provide fuller palettes. My work here seeks to do both: to measure an

untapped signal continuously for long term evaluation. First, it may help to see where this signal, intravascular volume status, finds itself among the current state of the art.

4.1.1.1 Vital signs, vital signals

Among the many medical measurements at our disposal are the so called vital signs, metrics of essential body functions. The four standard vital signs are temperature, heart rate, respiratory rate, and blood pressure. These are measured in myriad ways— through thermistors, electrodes, advanced optics, piezoelectrics, sphygmomanometers—and, thanks to the rise of low power, low profile electronics, in myriad places. A large and growing industry of commercially available devices for hospitals and consumers able to track one or more of these signals has sprung up in the past decade due to a confluence of factors. On the medical side the positive impact of outpatient care couple with the widespread availability of telecommunication has enabled patients to return home earlier, while monitored more closely after initial discharge from a care center (when many complications from procedures or missed problems first arise). This has been matched on the consumer side by the growth of quantified self initiatives and the Internet of Things where the desire to know more about ourselves and to have our products act upon that knowledge has spawned the rapid growth of many industries just across my short graduate school tenure.

Of those advanced wearable health monitoring devices, few measure more than two of these signals (usually temperature and heart rate). The respiratory rate will likely be the next frontier for fitness trackers. Since most of these trackers already use photoplethysmography to track the heart rate and since respiration can be detected by such means (recall Figure 1.13 in Chapter 1)¹, this would be an obvious next step. In fact, I suspect there are researchers in some companies working on this problem as this dissertation goes to print. A review of blood pressure monitors, however, quickly reveals that a technique beyond the sphygmomanometric or the piezoelectric is needed for to make a wearable, non-obstructive monitor of this signal a reality. Again, I suspect a photoplethysmographic approach of such a measure (via waveform analysis) will be the manifestation we see in clinical and commercial practice. Moreover, I would imagine a company that could accurately measure all four in a wearable and durable package will find a large market for its wares in both consumers and health systems.

¹One may also detect respiration via its effects in varying the heart rate.

4.1.1.2 The fifth vital

To these four is often added a “fifth vital” to give a fuller portrait of a patient’s prognosis.² Depending on the context of implementation, such supplementary vitals have included oxygen saturation [216], pain [217], and spirometric pulmonary function [218]. Each of these methods, though useful, are useful only to an extent. Taking these signs in reverse order, we might be immediately struck by the obvious limitation of pulmonary function as a vital sign: one must breathe to make it. Though the opinion paper suggesting its use believed that the instantaneous assessment of airflow obstruction could aid in clinical decision making, little to no research has shown its influence on patient management or improved patient outcomes [219]. Though many of individual parameters of spirometric testing helpfully assess many conditions (asthma, chronic obstructive pulmonary disease, cystic and pulmonary fibrosis, etc.), spirometry itself as an indicator of whole body functioning at the level of the other four vital signs is not sufficient.

Pain is perhaps most consistently cited as the “fifth” vital. Spearheaded by the Department of Veterans Affairs in the 1990s [220], pain management as a guide for therapy has been implemented internationally for over decade now. With roots entrenched from the beginning of modern medicine (“First do no harm.”), the impetus to treat pain as a vital sign stems from the belief that physicians act on vitals and would therefore be more inclined to treat and manage pain more effectively if elevated in status. Unfortunately, there is evidence that suggests the measurement of acute and chronic pain does not improve patient outcomes nor improve the quality of pain management [221, 222]. The subjectivity of pain—its multiple facets, its complex manifestation, our poor articulation—make acting on its presence and absence difficult. Without quantification, pain as a signal whose progress physicians can track and modify to guide treatment is uncertain.

Oxygen saturation has also seen considerable support as the fifth vital [216] and is perhaps most deserving of the title. Quantifiable, measurable, non-invasive, and objective, it satisfies many conditions that pain does not. For all those benefits, though, it still has a few detractors. First of all, oxygen saturation is only a surrogate for the physiologically more meaningful arterial partial pressure of oxygen (PO_2). Nor can it estimate any other blood gases (such as PCO_2) or pH. Moreover, in healthy subjects oxygen saturation is nearly always at or near 99%, and thus may not attract the attention of consumers interested in monitoring their own health.³

²Researchers have thus far branded sixth and even seventh vital signs, though they have not seen mainstream adoption.

³As silly as it may sound, dynamically varying data likely spurs more interest than invariant data. After all, there is a reason we do not measure our heights from day to day, even during our fastest growth spurts.

Other medical signals that are useful though not vital and commonly recorded in specific situations include actigraphy, blood glucose, electrocardiography, electroencephalography, electromyography, and galvanic skin response. Each has a wearable manifestation on the market with a varying degree of accuracy and success. A list of these signals, including my own impedimetric intravascular volume evaluation, can be found in Table 4.1.

I would like to propose intravascular volume status as another in the field of fifth vitals. It is an independent, actionable, and measurable signal that reflects the overall well being of a person. It is vital to health and reflective of disease. If any sign deserves its place among the vital parthenon, may it be this.

4.1.2 The potential role of bioimpedance

Bioimpedance, as shown in the previous chapter, is uniquely suited to easily measure a few of the vitals, including my proposed new vital, in a highly detailed fashion. High quality reconstruction of the arterial waveform easily tracks the heart rate while the respiratory rate (and magnitude) is seen in the venous shifts of blood. There is even evidence to suggest the arterial waveform itself is indicative of blood pressure [223, 224, 225], and would thus leave only temperature to be added. Taking advantage of both the electrical and the chemical, the anatomical and the physiological, bioimpedance in the right context can reveal far more about the inner workings of a person than many other comparable modalities. Thus, I contend that bioimpedance ought to play a role in the future of wearable biomedical monitoring.

Limb impedance satisfies what I consider to be the three critical conditions for a continuously monitored signal: it is universal, unobtrusive, and meaningful. By universal, I mean that everyone has blood and cells, tissues and fluids, enough ionic solutions and dielectric materials throughout their body such that everyone is able to be measured. Furthermore, those measurements can be made without obstructing the habits of the user. All that is required are a few points of electrical contact, which may be accomplished through adhesive, dry, or textile electrodes. Finally, the results it yields are meaningful in that they reflect current bodily functions, suggest clinical options, and can track the progress changes (as was the case in the dialysis study of Chapter 2). With each of these conditions fulfilled, I feel justified in having pursued a wearable device to continuously monitor the limb impedance of a person, as detailed in the sections that follow.

Table 4.1: A list of a few biomedical signals currently able to be measured through wearable electronics. Listed are the signal, a few sensors types currently used to measure them, and the signal source. Dividing the five vitals at the top of the list and several other practical signals is my own I2VE, whose status I believe ought to be treated as a measure of a vital sign (volume status) whose time shall come.

Biosignal	Sensor type	Signal source
Temperature	Temperature probe, skin patch	Body's metabolism warming tissue through energy expenditure
Heart rate	Electrodes, phonocardiograph, pulse oximetry	Variations in parameters caused by the cardiac cycle (electrical dipoles, pressure, volume, etc.)
Respiratory rate	Piezoelectric, piezoresistive	Chest motions corresponding to breathing
Blood pressure	Sphygmomanometer	Pressure exerted on the vessel walls by circulating blood
Oxygen saturation	Oximeter (optical)	Differential optical absorption characteristics of bound and unbound hemoglobin
I2VE	Electrodes	Respiratory induced variations of peripheral blood volume
Actigraphy	Accelerometer, gyroscope, inertial motion unit	Motion, position, or orientation of the body or its parts
Blood glucose	Glucose meter	Concentration of sugar within the body regulated by metabolic homeostasis
Electrocardiography	Electrodes	Movement of the heart's electrical dipole
Electroencephalography	Electrodes	Electrical potentials generated by brain activity
Electromyography	Electrodes	Muscle activation induced electrical phenomenon
Galvanic skin response	Electrodes	Electrical conductance of the skin as a function of sweat gland activation

4.2 A wearable bioimpedance monitor

To extend the results of the physiological investigations presented in Chapter 3 to more environments, I attempted to design and create a small, low-profile device that could measure bioimpedance as accurately as a bedside monitor.⁴ Put simply, such a wearable device had merely two functions: to measure the impedance of an object and to communicate that result. Priorities were placed on ensuring that this acquisition and communication could be done in real time and in as many different scenarios as possible. As such, we sought to develop a device capable of rapidly measuring impedance with different electrode configurations (both bipolar and tetrapolar), different electrode positions (wrist, arm, leg, etc.), and communicate via a wired or wireless protocol. Furthermore, we wished to do so with more mobility than current benchtop and bedside monitors afford. Hence, we sought to design a wearable research platform for monitoring bioimpedance in as many places and in as many ways as someone could hope for.

In fabricating such a wearable, I hope to join those in the philosophy that physiology does not happen within the hospital or the home but within the patient, where ever they may come.

Figure 4.1 presents a simple block diagram outlining the high level operation of the I2VE wearable. Again, because the point of this device is merely to measure the impedance of an object and transmit that data for researchers to review, two dominant wings flank the processor in this figure. The first (on the left) demonstrates the impedance measurement work flow. Electrodes placed on the object to be measured are connected via wires to the body of the device. Once there, a custom designed reconfigurable frontend paired with a network analyzer (AD5933) excites the material to elicit an impedance response. This impedance response data is then converted to a format that the processor can work with. Once through the processor, the data enters the second wing (on the right) to communicate that data to a user, a researcher, a physicians, etc. A Bluetooth radio facilitates wireless communications while a micro-USB (universal serial bus) enables wired data transmission, programming, and battery charging. (The reception of this data by a receiver is not shown nor discussed at any great length here. This is because any device capable of receiving Bluetooth communications can be pressed into service.⁵)

⁴I was not alone in this endeavor. Adetunji Dahunsi, Il Taek Kwon, and the people at Henway Technologies all contributed significantly to the work presented here. Though I am giving voice to the work done, these individuals was invaluable in their own ways in the design, creation, and validation of this device. This final chapter would not have happened without them and I hope it stands as a testament to our combined efforts. I have used the plural first-person throughout to further emphasize this collective effort.

⁵In part its absence from this dissertation is due to Dahunsi's excellent work on this front, which always made the process seem like magic, though a great deal of engineering went into it.

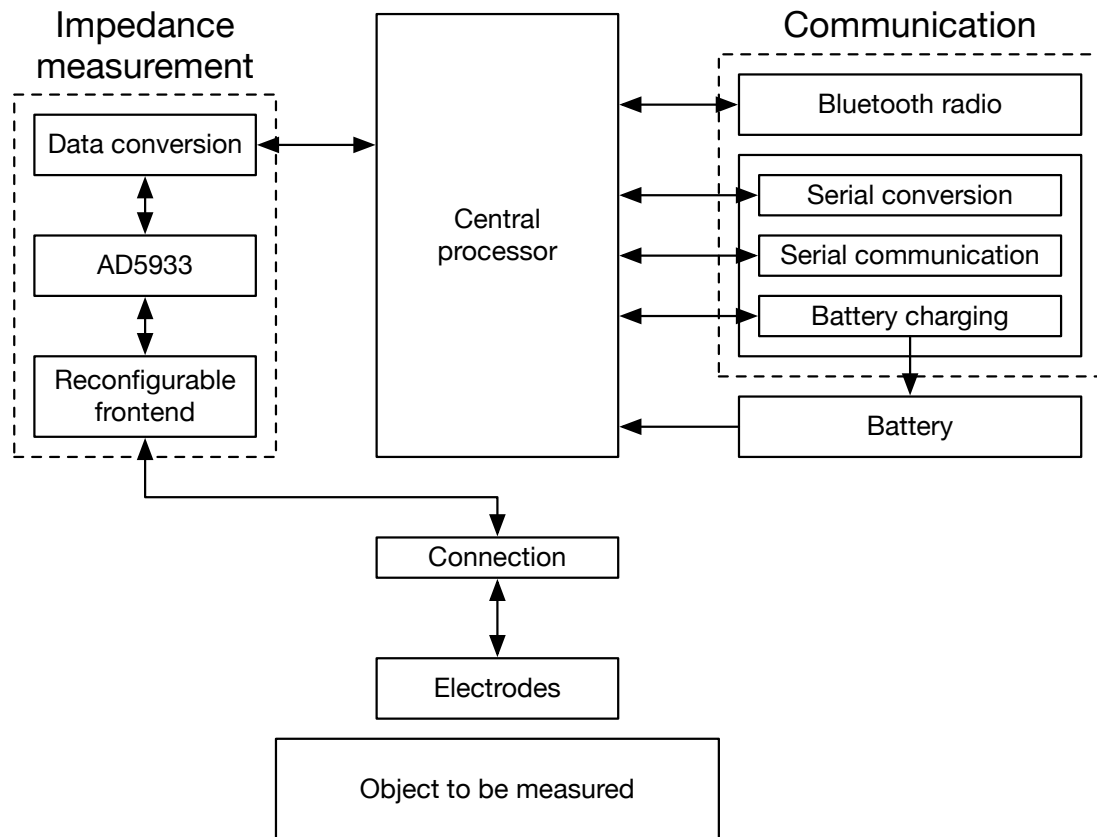


Figure 4.1: A block diagram representation of the wearable bioimpedance monitor developed. Electrodes are placed on an object and connected via a 3.5 mm audio jack to the device. The device itself then measures the impedance, sends that data to a central processor, then transmits that data either wirelessly (Bluetooth) or through a USB connection. The USB connection also serves to program the device and charge the battery. Each of step and component of this process is discussed in this chapter. The actual printed circuit board based upon this diagram used in this work can be seen in Figure 4.2(b).

This approach in tow, a circuit was designed, tested, manufactured, and validated. A simplified representation of the printed circuit board and the physical board of the wearable device itself are shown in Figure 4.2. In this figure, the minor components (such as switches, battery connectors, lights, and a prototyping area) can be viewed in relation to the necessary components (the 3.5 mm audio jack connection for the electrode wire, the micro-USB connector, the Bluetooth antenna, the AD5933, and the microprocessor), to give a fuller picture of what all went into this device. On the underside (not shown) are passive elements only (resistors, capacitors). Experts in the miniaturization of electronics will note that by not utilizing the underside of the device to carry more components and by not minimizing the distances between components that this device is not nearly as small as it could be. This is correct. What I present here should be considered a demonstration of the concept of wearable bioimpedance monitoring and not a prototype ready for full scale deployment.⁶

What follows is a discussion of the design of this device and its constitutive components, its validation both *in vitro* and *in vivo*, and what I believe can be done to improve this and other wearable medical monitors in the future.

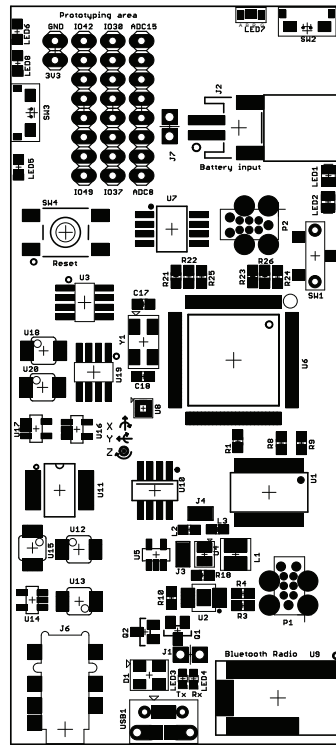
4.2.1 The design

Reviewing the literature on the subject, many researchers have turned to the AD5933 impedance analyzer chip to lie at the heart of their bioimpedance monitors [226, 227, 228, 229]. Our design follows their lead. Produced by Analog Devices, the AD5933 is a high precision (12-bit) impedance converter system that works by exciting a material between two terminal at a known potential and measuring the induced current flow [230].⁷ This is distinct from other approaches (including other integrated chip approaches, such as Texas Instruments AFE4300 [231]) where current is applied to an object and the potential generated is measured to determine the impedance. An disadvantage of the former approach (voltage excitation) that the latter (current injection) solves is that one may not know the current one is applying and thus could possibly send in a large amount that could damage a material or harm a subject. However, controllable current sources are notoriously fickle and the parasitic capacitances of electronic circuits can make it difficult to achieve desired values without tuning. Weighing these options, we decided to use the AD5933 in our own work, modifying it as necessary to measure our impedances of interest.


Once an instantaneous impedance of a subject is measured, that data must be processed

⁶As of the time of this writing, I and my collaborators have worked with Henway Technologies on a revised version of this device with a footprint smaller than a standard credit card.

⁷Recall from equation ?? that the ratio of these two parameters (voltage and current) give impedance.



(a)



(b)

Figure 4.2: A (a) simplified representation of the printed circuit board and (b) the actual printed circuit board itself. In the top left corner is a button flanked, in the top right is another button and LED combination, the battery connector, and the sliding switch to turn the device on and off. Towards the middle of the board is a reset button, microcontroller, and AD5933 and its analog front end. At the bottom from left to right is a 3.5 mm audio jack connection, a micro-USB connector, and the Bluetooth antenna. (Actual size if printed on 8.5" x 11" paper.)

and communicated outward. These three main design thrusts—to measure impedance, process the data, and communicate it—underlie all the work that follows. That is, every decision chosen and every experiment conducted were done to either prove or improve the function of the wearable device. The specific implementations of each of these thrusts are outlined below.

4.2.1.1 Measurement: the AD5933 and its extension

The AD5933 works by exciting a material with a known voltage at a known frequency, measuring the current flowing through the material, and computing the complex impedance via their ratio. More specifically, the AD5933 network analyzer operates by feeding a numerically synthesized sinusoidal signal generated at a specified frequency (between 1 and 100 kHz) by a direct digital synthesizer core to a digital to analog convertor which is then passed through to an adjustable output gain amplifier. As the material is excited by the induced potential, a current-to-voltage converting amplifier (with a gain adjustable via the current converting feedback resistor, RFB) detects the current response and feeds it through a programmable gain amplifier, an anti-aliasing filter, then finally through a 12-bit analog-to-digital converter. This digital signal then undergoes a discrete Fourier transform via a digital signal processing engine to estimate the spectral power of the current excited by the voltage at the selected frequency. The ratio of this current measurement with the known potential input and gain then gives the real and imaginary components of impedance, which are fed to an inter-integrated circuit (I2C) communication interface to be transmitted outward for data analysis and storage. A block diagram showing this operation can be seen in Figure 4.3.

Though the AD5933 has been used successfully by many researchers, some have noted that its bipolar configuration

Though there are three stages—voltage excitation, current measurement, and ratio determination—compose

are labeled in the general block diagram representation of the AD5933 shown in Figure 3.

While the AD5933 is an excellent chip and a real boon to bioimpedance research, it is not without its shortcomings. For my work, three key limitations were present: a fixed range of potentials, a positive DC bias, and a native bipolar configuration, which I have expounded upon below.

1. **The range of output excitation voltage amplitudes is from 198 mV to 1.98 V peak to peak.** Though these levels are appropriate for the macroscopic biology of

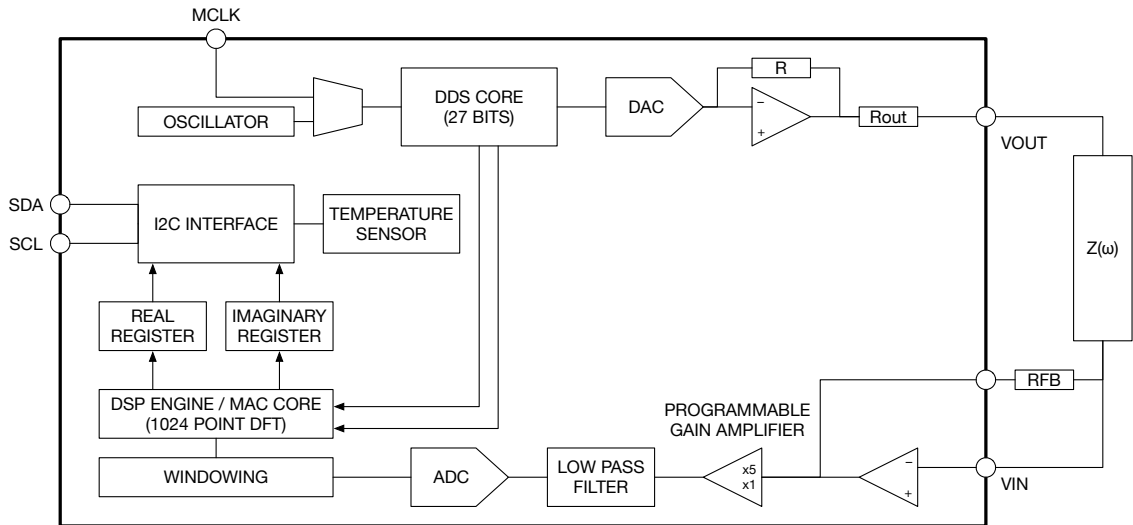


Figure 4.3: Operation of the AD5933. A sinusoidal voltage signal at a designated frequency is generated via the combination of the master clock (MCLK), an oscillator, and a direct digital synthesizer (DDS). This discrete signal is then converted to analog and amplified, reaching the material as VOUT. Impedance in the material, $Z(\omega)$, is found by passing the current at VIN through an amplifier, filtering, converting to a digital signal, then taking the discrete Fourier transform at the digital signal processing (DSP) engine with information supplied by the DDS core. The real and imaginary components of the impedance are then passed to an inter-integrated circuit (I2C) communication interface and transmitted out through serial data (SDA) and serial clock (SCL) lines.

human beings under most circumstances, for very large or very small objects these potentials could pose the problem of either themselves being too small or too large, respectively. Since we are dealing with problems of miniaturization here, let us put to one side the scenario of measuring large biological materials. Reducing the size of the electrodes increases current density, which can negatively impact the material, the electrodes, and interface between the two.

2. **The excitation signal output varies around a DC bias for these ranges from 173 mV to 1.48 V.** Again, while acceptable in numerous situations for which the chip was designed, in cases of prolonged measurement, electrochemical processes at the electrode-electrolyte interface may alter the object one is intended to measure. (This is especially relevant for metal electrodes interfaced with an aqueous solution.)
3. **The chip is design for a bipolar configuration.** Perhaps the most damning indictment made of the chip from the bioimpedance community is its bipolar configuration known to introduce often significant electrode-electrolyte interface effects (usually reported under the banner of “electrode polarization” [232, 233]). In many situations

these effects can dwarf those that one intends to measure. As such a bipolar measurement may only report the high impedance of the interface, remaining insensitive to impedance of the object one wishes to monitor. Moreover, the density of current near the electrodes varies widely, proximally affecting the region non-uniformly.

Each of these limitations suggests their own improvements. The fixed range of output potentials can be corrected by the addition of a gain modifying circuit. In this case, where we seek to reduce the lowest possible potential to minimize current density at the electrodes, we could implement an attenuation circuit. To remove the DC bias, a simple high-pass filter can be applied. Moreover, a voltage follower to couple the AC signal could be used to buffer and track the signal more robustly. To mitigate the effects of a bipolar setup, we could redesign it to be tetrapolar. Separating the current carrying and voltage sensing electrodes significantly reduces the artifacts of polarization because the only set of electrodes affected are not the same as those measuring the impedance. To press the AD5933 into such service, one could convert its output excitation potential, VOUT into a driving current, sense the potential through another means, and convert the sensed potential back to current to be measured at VIN. This was the approach of Seoane et. al [226] and was used to great effect in their work. However, as they and others have noted, the current source can become unstable if used over a large range of excitation frequencies and impedances. To properly tune the current source would require a more complex approach. Instead of attempting to drive current, Margo et al [229], whose work we follow closely here, instead continued to utilize voltage excitation and simply measure the current and voltage sequentially in a two step process. Their reported results were incredibly accurate and precise over a large range of frequencies and impedance values. As such, it was the starting point of our own design.⁸

We designed a reconfigurable analog front end to extend the functionality of the AD5933, solve the problems outlined above, and give researchers the option to use either a bipolar or a tetrapolar setup based on the circumstances of deployment. Figure 4.4 shows a block diagram representation of the circuit used to modify the operation of the AD5933. Starting at VOUT, the excitation signal is attenuated, high-pass filtered, and AC coupled before being sent to the electrode lead connection point, which in our case was chosen to be a 3.5 mm audio jack to make the connection and removal of electrodes a one step procedure. The two inner electrodes (S+ and S-) sense the potential of the material via an instrumentation amplifier (an AD5226, a low power alternative to the high speed AD8250 used by Margo et al.) whose voltage value was converted to a current (via a traditional transimpedance am-

⁸Our ends differ markedly from those of Margo's and as such discrepancies will accrue here and there. Where Margo et al. wished to go microscopic, we wish to create a monitor for the masses.

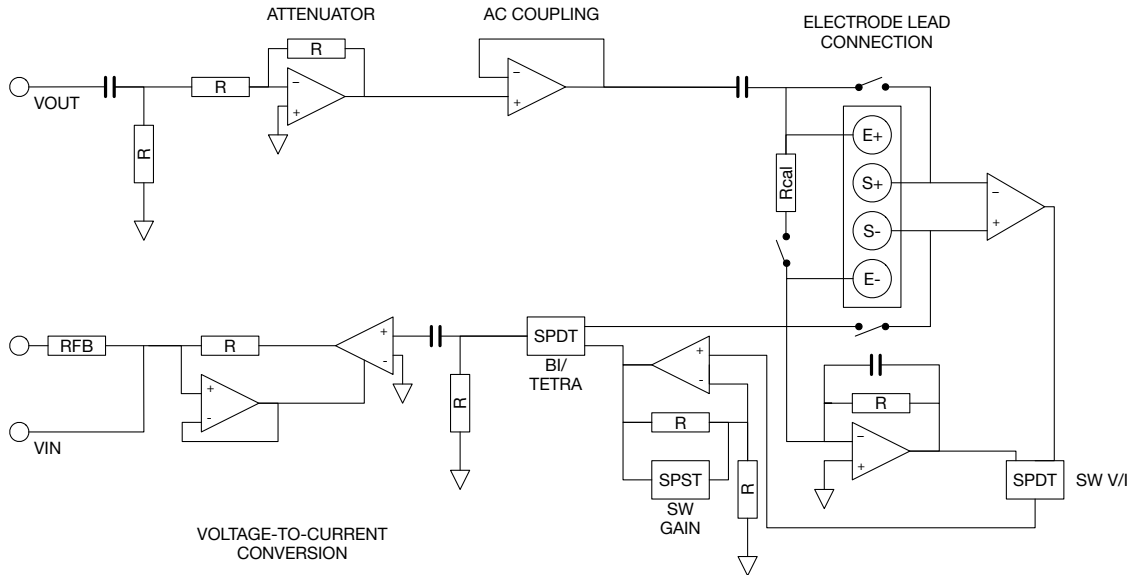


Figure 4.4: The reconfigurable frontend designed to work with the AD5933 shown in Figure 4.3. Such a front end allows a bipolar or tetrapolar electrode arrangement to be specified on the fly, enabling researchers to choose that which best fits their needs.

plifier) to be measured across the feedback resistor, RFB, and at VIN. Added to Margo's original work is the ability to switch between their tetrapolar setup and the bipolar setup native to the AD5933 itself. This software switch allows the setup to bypass the additional tetrapolar components and operate normally, albeit while having to compensate slightly for the added circuitry.

As such, we have developed the first extension circuit for the AD5933 that can switch modes on the fly, giving users the ability to either avail themselves of a tetrapolar arrangement if they should so fear interface or proximal effects in a situation or a bipolar arrangement if these are not a concern.

4.2.1.2 Processing: the microcontroller in context

The microcontroller chosen for this project was an ATmega1280 (from Atmel), an inexpensive, low-power 8-bit AVR RISC-based processor with 86 general purpose input/output lines, 32 general purpose registers, real and flexible timers, two-wire serial interface, a 16-channel 10-bit analog to digital converter, an operational range between 2.7-3.3 V, and a maximum throughput of 16 MHz. For the needs of this project, this processor may have been excessive, but was chosen to lay a firm groundwork for future iterations where more

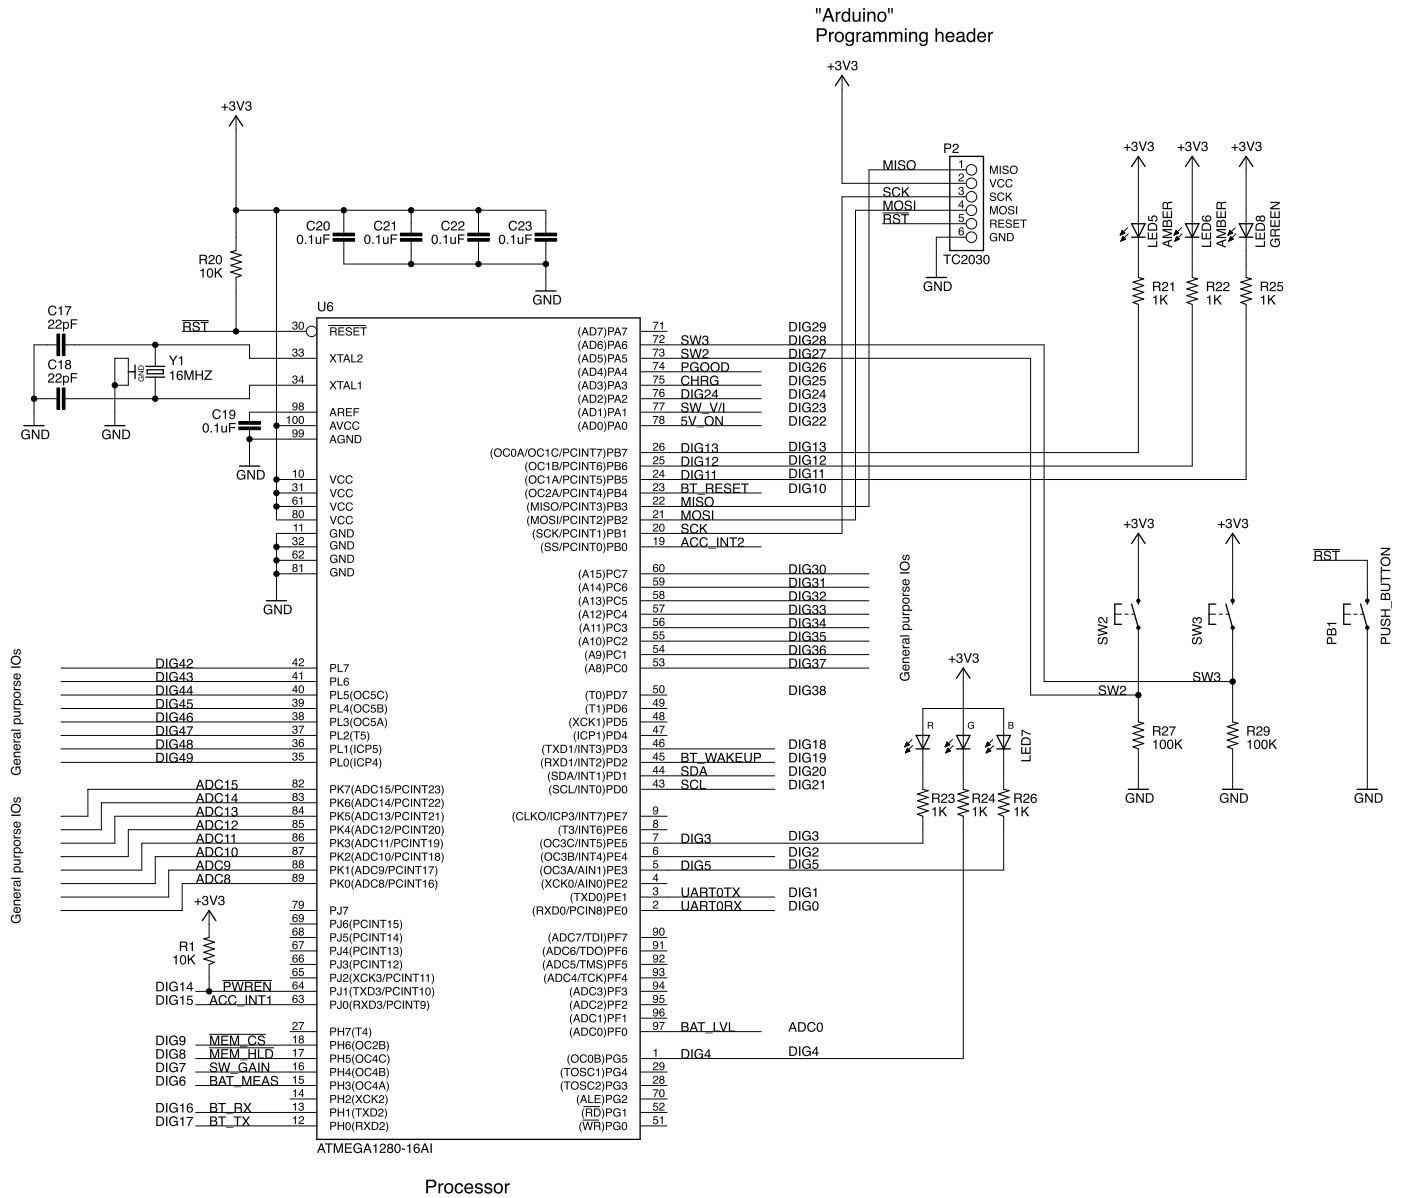


Figure 4.5: The processor (ATmega1280) of the wearable device and each of its relevant connections. Each of the main functions of the device has a presence here: communication with the AD5933 (PD0 and PD1), USB (PE0 and PE1), and Bluetooth (PH0 and PH1); memory storage (PH6 and PH7); and programming (such as PB1-3).

powerful processing may be necessary.⁹ Figure 4.5 shows each of the connections to the microcontroller from the various components of the device.

Starting in the upper left hand corner of the figure and going down we can see the reset function, $\overline{\text{RST}}$, created by a push button (aptly named `PUSH_BUTTON` on the far right of the figure), the clock generation (an oscillator connected between `XTAL1` and `XTAL2`), and a series of reference grounds and voltages (`AREF`, `AVCC`, `AGND`, `VCC`, `GND`), general purpose analog and digital inputs and outputs (`DIG42-49` and `ADC8-15`). Two interrupt signals are connected at Port J (specifically designed for this use). At `PJ1` is an input $\overline{\text{PWREN}}$ signal to interrupt operations when powering the device. Similarly at `PJ0` an interrupt signal from the accelerometer (a `BMA250`, discussed later), `ACC_INT1`, can be measured and acted upon as need.

Digital pins `DIG8` and `DIG9` (from port H, `PH6` and `PH5`) are used to supply the signals to select the memory chip (an `M25P32`) or to pause serial communications to the chip, $\overline{\text{MEM_CS}}$ and $\overline{\text{MEM_HLD}}$, respectively. From the same port (H): the signal to change the software gain of the reconfigurable front end (as seen in Figure 4.4), `SW_GAIN`, is read in and modified (from `PH4`); actionable battery measurements signals, `BAT_MEAS`, are reported (from `PH3`); and data received and transmitted by the bluetooth antenna (a `BLE113` from Bluegiga, discussed later) are communicated via `BT_RX` and `BT_TX`, respectively, from `PH1` and `PH0`, specifically designed into the microcontroller to receive transmit and receive data.

Along the right side of the figure starting at Port A, the uppermost signals fed into the processor are `SW3` and `SW2` stemming from their corresponding switches and allow one to program responses to one or both of the switches being engaged (such as pausing recordings, pushing data manually, taking new measurements, etc.). `PGOOD` and `CHRG`—signals generated by the battery charging circuit outlined in the next section—indicating if the device is charging are connected at `PA4` and `PA3`. The signal to throw the switch for the voltage and current multiplexer of the reconfigurable frontend, `SW_V/I`, comes from `PA1`. Finally, `PA0` provides the on-off signal, `5V_ON`, for the voltage converter (discussed in the next section)

Port B is an 8-bit bidirectional communication port with better driving capabilities than the others, and as such is equipped with serial peripheral interface (SPI) capabilities. Because of these properties it is used here to facilitate peripheral communications. To get a better sense of how this works, let us stray from the processor proper for a moment and examine such an interface that feeding into the processor in the positions `PB1-3`. This is

⁹The large community support and the direct link to the Arduino prototyping platform also factored into this decision.

a programming header (embedded in the board via a Tag-Connect TC2030) to allow developers to communicate with the board via a peripherally attached device such as another microcontroller, memory (more on that in a moment), other sensors, and to reprogram the microcontroller itself. This small, low-profile port was chosen to reduce the footprint required (minimizing the board), though it requires an additional plug if one wished to install drivers or reprogram the board (a TC2030-IDC-NL cable to a CC Debugger (from Texas Instruments) in our case). This header has six signals to work with half of which require no further explanation (a supplied voltage, VCC, a ground, GND, and the reset signal we have encountered previously, \overline{RST}). The other three are how communication between the header and the processor occur and should be familiar to those who have worked with serial peripheral interfaces before. They are the master-input-slave-output (MISO), the master-output-slave-input (MOSI), and serial clock (SCK), and follow the SPI protocol. Three digital pins (DIG11-13) control LEDs to indicate to a developer if and how the board is communicating with a peripheral (though, this can be reprogrammed however one would wish). From PB4 (DIG10) comes the Bluetooth reset command (BT_RESET). The final pin of Port B (PB0) takes in an accelerometer interrupt (ACC_INT2) similar to that found in Port J (ACC_INT1 at PJ0).

All of Port C is treated as general purpose digital inputs and outputs, though the special external memory functions available to the pins could be of use to future developers wishing to build in greater longterm analysis features.

Port D contains three very important pins. PD3 transmits the wakeup command of the Bluetooth antenna, BT_WAKEUP, to initialize the antenna for wireless communication. PD1 and PD0 are the serial data, SDA, and serial clock (SCL) lines, respectively. These two signals communicate with the AD5933 through the I2C (inter-integrated circuit) protocol. However, because the device itself is powered at 3.3 V (by a lithium polymer battery) and the AD5933 is set to a 5 V range (labeled as AVDD throughout this chapter), a translational step must be included. Such a step can be found in Figure 4.6 where the SDA and SCL of the microcontroller are converted by a dual bidirectional I2C voltage-level translator (PCA9306 from Texas Instruments) to SDA_5V and SCL_5V for the impedance measurements. Through these pins and through this translator all communications with the AD5933 must go.

Signals to and from the USB port (once converted, as discussed in the section that follows), UART0TX and UART0RX, are connected to pins PE1 and PE0, specified by the manufacturer for this purpose. The other pins of Port E are simply used to inform the developer of incoming and outgoing USB communications (via LEDs).

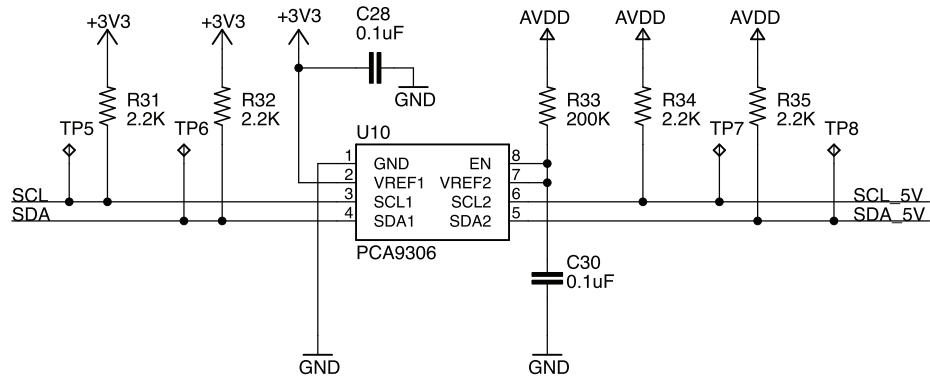


Figure 4.6: I2C voltage translator to make possible communication between the microcontroller powered at 3.3 V and the AD5933 powered at 5 V. Signals from the microcontroller, SDA and SCL, are stepped up to become SDA_5V and SCL_5V, to interact with the impedance analyzer.

4.2.1.3 Communication: USB and Bluetooth

There are two main ways with which to communicate with this device (as shown in Figure 4.1): USB and Bluetooth, a wired and wireless standard, respectively. As with many other decisions for this device and indeed this work, the multimodal approach was chosen to make it as accessible and extensible as possible. Researchers and developers have the option to work in the way and with the protocol they feel most comfortable. In our practice, programming over USB was considerably easier than to do through the TC2030 header outlined in the previous section, while transmitting data through Bluetooth made real-world measurements more straightforward than any wired transmission. There were, thus, benefits to utilizing both, though it slightly increased the complexity and size of the device.

USB connection: programming, communication, and battery charging

Starting from the micro-USB connection of Figure 4.7, one can see that we use four active pins: VCC providing +5 V, D- (Data-), D+ (Data+), and a ground (GND).¹⁰ To protect the device in general from voltage spikes arising from the USB port (usually during plugging and unplugging a cable), a transient-voltage-suppression circuit was placed between each of the inputs and ground (using a PRTR5V0U2x fro NXP Semiconductors). The VCC pin delivers a potential, V_USB, used through the device when connected. This is the potential used as the input (at pin 1) for the battery charging circuit, a BQ24090 chip from Texas Instruments. The output potential is then fed to a 3.3 V, 1000 mAh lithium polymer

¹⁰A fifth pin, ID, located between D+ and GND allows electrical identities to be given to cables. Since we make no use of it, we will pay it no more attention.

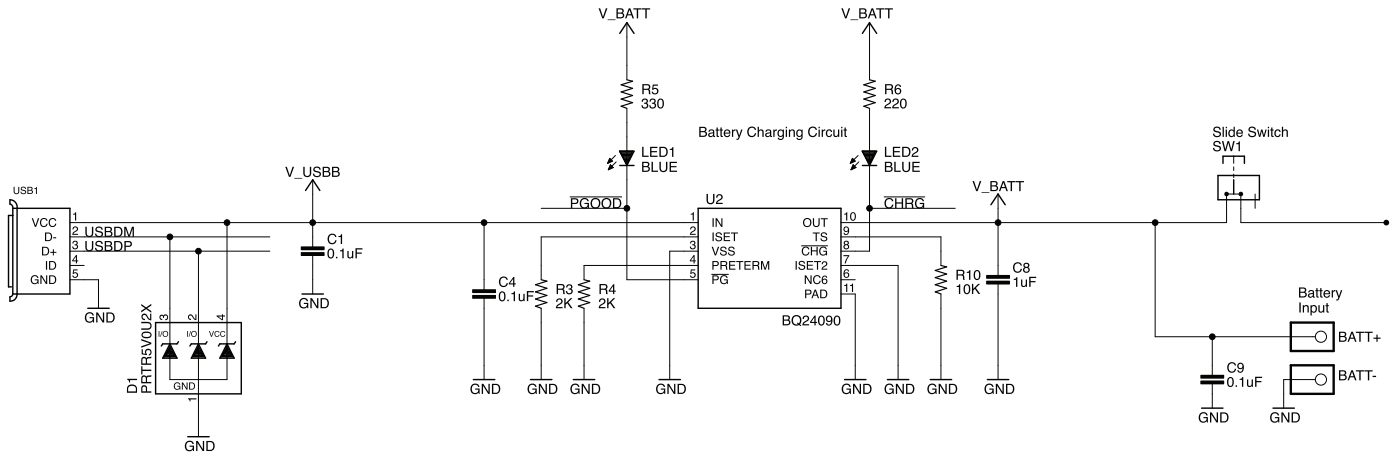


Figure 4.7: USB connection and the battery charging circuit. Potential from the USB port charges the battery which in turn powers the device. Signals from D- and D+ (USBDM and USBDP) are sent to the USB to serial converter seen in Figure 4.9.

battery. As mentioned previously, the status of the charging state is reported via the signals $\overline{\text{PGOOD}}$ and $\overline{\text{CHRG}}$ to the processor and made immediately clear to the user as a pair of LEDs.¹¹ Potential from the battery, V_BATT, is used to power the device. To turn the device on, the switch SW1 must be flipped to the correct position.

Since the device makes use of two potentials (3.3 and 5 V), once the switch is flipped to power the device, the potential is sent to both a linear regulator (a TLV70233 from Texas Instruments) and a voltage converter (a LTC3525-5 from Linear Technology), as shown in Figure 4.8. The regulator is one designed specifically to work with low quiescent current components and was thus ideal for the I2C translator (Figure 4.6), the USB to serial converter (Figure 4.9), the Bluetooth radio (Figure 4.10), the accelerometer (Figure 4.11), and memory (Figure 4.12). The voltage converter was chosen to step up the amplitude of the DC potential supplied by the battery to use the widest possible range of the AD5933 and its frontend, while remaining mindful of low current requirements (the LTC3525-5 has a quiescent current $\leq 7 \mu\text{A}$). This portion of the circuit ensured minimal current leakage and power drain, helping prolong battery life.

The two data signals from pins D- and D+, USBDM and USBDP, seen in Figure 4.7 are sent to a USB to serial UART (universal asynchronous receiver/transmitter) converter, as shown in Figure 4.9. This portion of the circuit is meant to simply translate the digital signals from a source utilizing the USB port (such as a computer) to serial data that can then be transmitted to and received from the rest of the device as the signals UART0TX and UART0RX. With this converter, data could be transmitted directly to a computer making

¹¹ $\overline{\text{PGOOD}}$ and $\overline{\text{CHRG}}$ are PGGOOD and CHRG through MOSFETs.

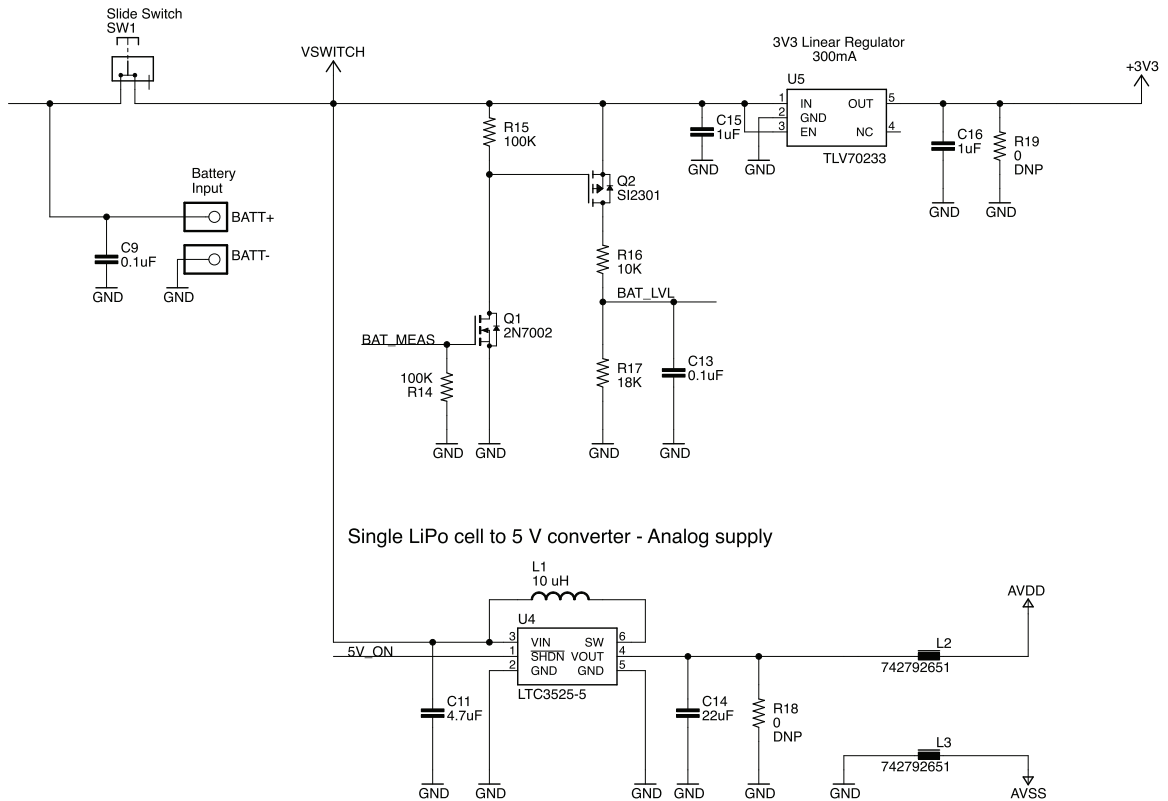


Figure 4.8: Battery regulator components of the circuit ensuring that consistent potentials of 3.3 and 5 V are available to those components of the device that require them. Both the linear regulator and the voltage converter are designed with the low quiescent currents of wearable electronics in mind.

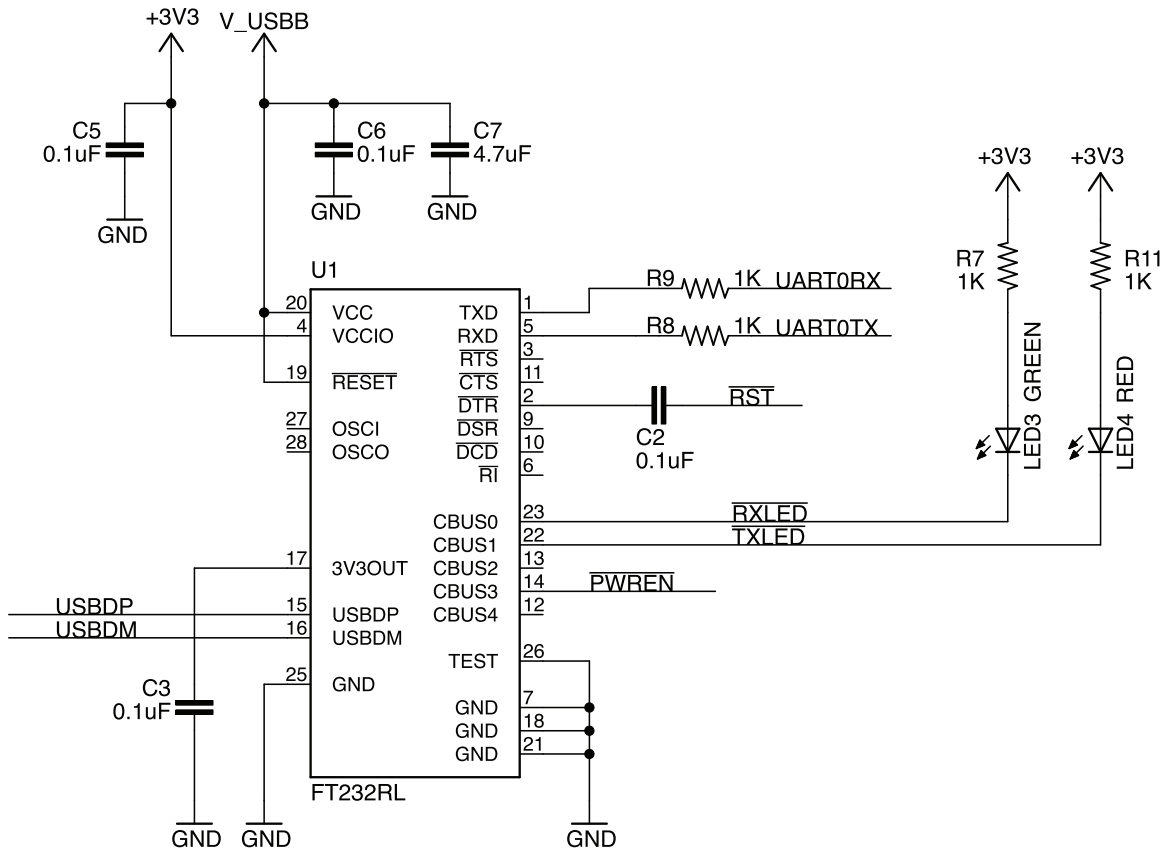


Figure 4.9: USB to serial converter that enables wired communication—be it programming, data transmission, or data reception—between the device and a computer.

possible rapid debugging, quick and easy programming, and aiding in more advanced real-time analysis .

Bluetooth: a mobile protocol

At this stage of development, the decision over which wireless communication protocol was used was ultimately one of preference, convenience, and chance. Depending who one asks and what references one consults any number of wireless protocols would have fit the bill. In the wearable electronics space, three major standards exist: Wi-Fi (IEEE 802.11(a,b,g,n)); Bluetooth (IEEE 802.15.1); and ZigBee (IEEE 802.15.4). Each comes with their own slew of pros and cons. We chose to work with Bluetooth 4.0 (Bluetooth low energy) because of the available documentation, ease of both hardware and software implementation, large commercial footprint, low power requirements, and sufficient transmission range. Cases can be and have been made for either of the other main wireless standards. Again, our decision was one of preference and time. Better minds will decide which standard ought to reign in the wearable medical device sphere in years to come.

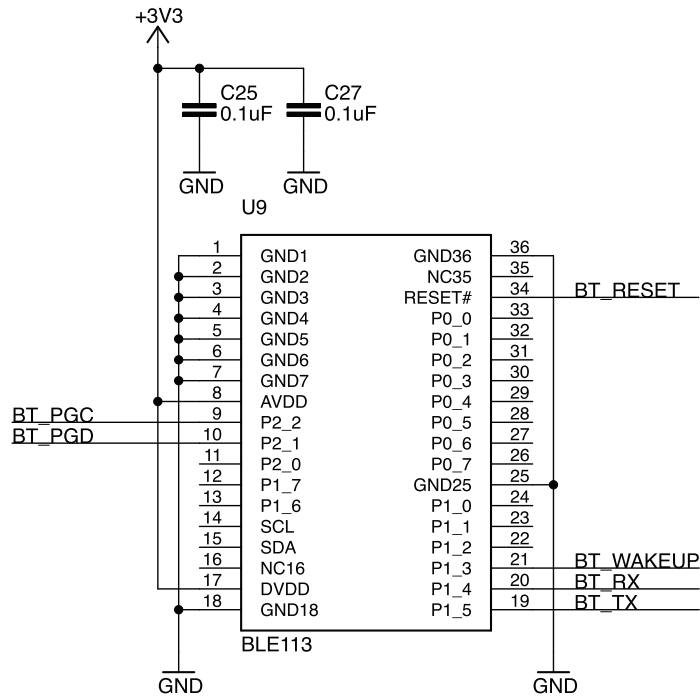


Figure 4.10: Bluetooth radio.

The radio chosen for this project (seen in Figure 4.10) was a BLE113 from Bluegiga Technologies, supporting Bluetooth low energy communication. For our work, six signals were of major importance: two to program (BT_PGC and BT_PGD); two to transmit and receive data (BT_TX and BT_RX); one to wake the radio (BT_WAKEUP); and one to reset the radio (BT_RESET). As we merely wished to transmit the data measured by the AD5933 and sent to the ATmega1280, much of the additional functionality of this radio is not utilized. Again, the goal of this device was to measure impedance and communicate that measurement. The singular use of many chips, including this Bluetooth radio ought to reflect this.

4.2.1.4 Miscellaneous: accelerometry, memory, and a list of components

We included an accelerometer for future actigraphy monitoring and motion compensation. We chose the BMA250 from Bosch as it is a very small (2x2 mm), low-profile (0.95 mm) triaxial accelerometer capable of tracking a wide range of accelerations (± 2 to ± 16 g) and patterns (tapping, orientation recognition, flat detection, etc.), with stand-alone and interruption capabilities (reported elsewhere), the schematic of which is in Figure 4.11. Communication to and from the processor with the accelerometer follows the same I2C communication protocol along the same SDA and SCL channels discussed previously.

Figure 4.12 shows the final component I will specifically call out for attention, a 32

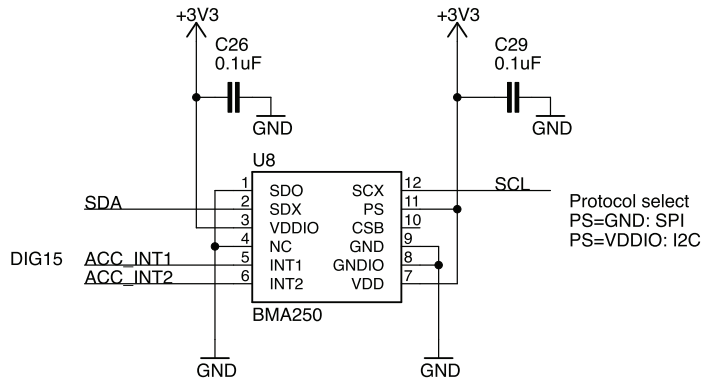


Figure 4.11: Accelerometer for actigraphy monitoring, motion activation, and compensation.




Figure 4.12: Local storage of measured data before transmission. Over two full days of complete impedance data can be collected continuously before it must either be deleted or transmitted.

Mbit EEPROM to store data locally before transmitting it wirelessly. Each of the signals here ought to be familiar as we have encountered them previously: $\overline{\text{MEM_CS}}$ and $\overline{\text{MEM_HLD}}$ control the memory states while MISO, MOSI, and SCK communicate with the chip directly via the SPI protocol.

Each of the components used in this device (and their in-circuit identifiers, values, manufacturers, and part numbers) can be read from Table 4.2.

Table 4.2: List of components used for the wearable electronic device. (Unless otherwise stated all passive surface-mounted devices are 0603 package sizes.)

Part Reference	Value	Manufacturer part number	Manufacturer	Description
C1,C2,C3,C4,C5,C6,C9,C10, C12,C13,C19,C20,C21,C22,C23, C24,C25,C26,C27,C28,C29,C30, C31,C32,C33,C34,C35,C36,C37, C39,C40,C43,C44,C46	0.1uF	GRM188R71E104KA01D	Murata	CAP CER 0.1UF 25V 10% X7R
C14	22uF	EMK212BBJ226MG-T	Taiyo Yuden	CAP CER 22UF 16V 20% X5R
C17,C18	22pF	C0603C220J5GACTU	Kemet	CAP CER 22PF 50V 5% NP0
C42	33pF	C0603C330J5GACTU	Kemet	CAP CER 33PF 50V 5% NP0
C7,C11	4.7uF	GRM188F51A475ZE20D	Murata	CAP CER 4.7UF 10V Y5V
C8,C15,C16,C38,C45,C41	1uF	GRM188R61E105KA12D	Murata	CAP CER 1UF 25V 10% X5R
D1	N/A	PRTR5V0U2X-215	NXP	TVS DIODE 5.5VWM SOT143B
J2	N/A	S2B-PH-SM4-TB(LF)(SN)	JST	CONN HEADER PH SIDE 2POS
J6	N/A	SJ-43617-SMT	CUI	CONN JACK 3.5MM R/A 4POS MID
L1	10uH	7440329100	Würth Electronics	FIXED IND 10UH 450MA 390 MΩ
L2,L3	600 Ω	742792651	Würth Electronics	FERRITE BEAD 600 Ω 1.0A
LED1,LED2	Blue	LB Q39G-L2N2-35-1	Oscram	LED CHIPLED BLUE 470NM
LED3	GREEN	LTST-C190GKT	Lite-On	LED GREEN CLEAR
LED4	RED	LTST-C190CKT	Lite-On	LED RED CLEAR
LED5,LED6	Amber	LTST-S270AKT	Lite-On	LED AMBER CLEAR RT ANG
LED7	RGB	MSL0201RGBW1	Rohm	LED CHIP RGB WHITE DIFF

LED8	Green	LTST-S270KGKT	Lite-On	LED GREEN CLEAR RT ANG
Q1	2N7002	2N7002	Fairchild	N-Channel Mosfet
Q2	SI2301	SI2301-TP	Micro	MOSFET P-CH 20V 2.8A SOT-23
R1,R10,R16,R20,R37,R40, R45,R46,R47,R48,R50	10 k Ω	ERJ-3EKF1002V	Panasonic	RES 10 k Ω 1/10W 1% SMD
R12,R13,R14,R15,R27,R28, R29,R30,R38,R39	100 k Ω	ERJ-3EKF1003V	Panasonic	RES 100 k Ω 1/10W 1% SMD
R17	18 k Ω	RC0603FR-0718KL	Samsung	RES 18 k Ω 1/10W 1% SMD
R18,R19,R41,R42,R49	0 Ω	ERJ-3GEY0R00V	Panasonic	RES SMD 0.0 Ω JUMPER 1/10W
R3,R4	2.00 k Ω	ERJ-3EKF2001V	Panasonic	RES 2K Ω 1/10W 1% SMD
R31,R32,R34,R35	2.2 k Ω	MCR03ERTF2201	Rohm	RES 2.2K Ω 1/10W 1% SMD
R33	200 k Ω	ERJ-3EKF2003V	Panasonic	RES SMD 200 k Ω 1% 1/10W
R43	511	ERJ-3EKF5110V	Panasonic	RES 511 Ω 1/10W 1% SMD
R44	12.4k Ω	RC1608F1242CS	Samsung	RES SMD 12.4K Ω 1% 1/10W
R5	330	ERJ-3EKF3300V	Panasonic	RES 330 Ω 1/10W 1%
R6	220 Ω	MCR03ERTF2200	Rohm	RES 220 Ω 1/10W 1% SMD
R7,R8,R9,R11,R21,R22, R23,R24,R25,R26,R36	1 k Ω	RC0603FR-071KL	Yageo	RES 1 k Ω 1/10W 1% SMD
SW1	N/A	CL-SB-12B-01T	Copal Electronics	SWITCH SLIDE SPDT
SW2,SW3	N/A	TL4100AF240QG	E-Switch	SWITCH TACTILE SPST
SW4	N/A	FSM6JSMA	TE Connectivity	SWITCH TACTILE SPST
U1	N/A	FT232RL	FTDI	IC USB FS SERIAL UART
U10	N/A	PCA9306D,118	NXP	IC LEVEL TRANSLATOR 8-SOIC
U11	N/A	AD5933YRSZ	Analog Devices	NETWORK ANALYZER 12B 1MSP

U12	N/A	LMV358IDGKR	Texas Instruments	IC OPAMP GP 1MHZ RRO 8VSSOP
U13,U18	N/A	AD8226ARMZ-R7	Analog Devices	IC OPAMP INSTR 1.5MHZ
U14,U17	N/A	AD8605ARTZ-REEL	Analog Devices	IC OPAMP GP 10MHZ RRO
U15,U20	N/A	TS12A12511DGKR	Texas Instruments	IC SWITCH SGL SPDT ANLG
U16	N/A	TS12A4516DBVR	Texas Instruments	IC SWITCH SPSTANLGL
U2	N/A	BQ24090DGGQT	TI	IC BATT CHRGR LI-ION 10MSOP
U3,U19	N/A	OP1177ARZ	Analog Devices	IC OPAMP GP 1.3MHZ
U4	5V	LTC3525ESC6-5#TRMPBF	Linear Technology	IC REG BOOST SYNC 5V 0.4A
U5	N/A	TLV70233DBVR	Texas Instruments	IC REG LDO 3.3V 0.3A
U6	N/A	ATMEGA1280-16AUR	Atmel	IC MCU 8BIT 128KB FLASH
U7	N/A	32-VMW6TG TR	Micron Technology	IC FLASH 32 MBIT 75 MHZ
U8	N/A	BMA250	Bosch	3-AXIS ACCELEROMETER
U9	BLE113	BLE113-A-V1	bluegiga	CLASS 2 BLTOOTH 4.0 MODULE
USB1	N/A	10118194-0001LF	FCI	CONN USB MICRO B RECPT
Y1	16 MHz	ABM3B-16.000MHZ-B2-T	Abracon	CRYSTAL 16MHZ 18PF




Figure 4.13: The calibration circuit used for the laboratory validation of the wearable device, reminiscent of those shown in Figure 3.5. Two potentiometers, R_1 and R_2 , and a capacitor, C , are meant to act as the extracellular and intracellular components.

4.2.2 The validation

To test how accurate this wearable device was both *in vitro* and *in vivo* tests were conducted. The former was a benchtop test on circuits meant to simulate the electrical equivalence of biological tissue, while the latter was conducted on six human volunteers. In both cases, outlined below, we saw good agreement between a standard measure and the device we created, making us hopeful that bioimpedance measures of volume status in real time and in the real world are now achievable.

4.2.2.1 Laboratory validation

The device was calibrated by connecting our device to an equivalent circuit made of one resistor, R_1 , in parallel with a series combination of a resistor and a capacitor, R_2 and C , respectively (in the form presented in Figure 4.13). This setup was composed of two varying digital potentiometers (AD5928, Analog Devices) and a set of four capacitors that we switched manually. The values of each of the digital potentiometers and the capacitors were measured using an LCR meter (Agilent E4980A) and for this experiment ranged from 356.4-1183.5 Ω for R_1 , 351.1-1146.4 Ω for R_2 , and 2.17-7.08 nF for C . To establish the devices accuracy and precision across a broad range of frequencies, resistances, and reactances, a frequency sweep was performed across each resistor and capacitor configuration. For the bipolar configuration the frequency was swept from 2-100 kHz and for the tetrapolar configuration the frequency was swept from 6-100 kHz. At frequencies below the minimum values (2 and 6 kHz for the bipolar and tetrapolar configurations, respectively), measurement was unstable.

To calibrate, the value measured at each frequency, resistance, and reactance for each configuration was then subtracted from its predicted value to yield the error in the mea-

surement. Two functional relationships were then fit between each of the parameters with respect to the measured value: a first order calibration curve and a second order calibration curve. The first order calibration curve sought to fit correct for a gain and offset in the measurement (either of resistance or reactance) in the form

$$R_{corrected} = R_{meas} + G_{cal}R_{meas} + O_{cal} \quad (4.1)$$

where, $R_{corrected}$ is the corrected measurement, R_{meas} is the original measured value, and G_{cal} and O_{cal} are the calibration curve gain and offset, respectively. The second order calibration curve extends the logic of the first order calibration curve and attempts to more accurately fit the shape of the errors likely caused by a complex interaction of many of the electronic components, such as the AD5933, multiplexers, digital potentiometers, etc. The second order calibration curve was of the form

$$R_{corrected} = R_{meas} + G_{1,cal}R_{meas}^2 + G_{2,cal}R_{meas} + O_{cal} \quad (4.2)$$

where $G_{1,cal}$ and $G_{2,cal}$ represent the fit gain parameters. From these calibration relationships, errors could then be removed in post-processing. One could do this on the device itself, but for the experiments reported here all calibration and correction was done after data collection. A further discussion on possible means of calibration, while fruitful, was beyond the scope of this study.

4.2.2.2 Human subject validation

Six spontaneously breathing volunteers were recruited for pre-clinical validation of the reconfigurable monitor. Subjects were asked to sit in a chair with their arms resting to their side on arm rests, breathing normally, while the bioimpedance of an upper limb was measured simultaneously by our calibrated monitor exciting at 50 kHz and sampling at 50 samples/second and a commercially available device, a Biopac EBI100C bioimpedance amplifier in conjunction with an MP150 data acquisition system (Biopac Inc. Goleta, CA), exciting at 50 kHz and sampling at 200 samples/second. Four electrodes were used for both devices (two outer exciting electrodes, two inner sensing electrodes), though when our front end was configured for a bipolar arrangement, only the inner two electrodes were used.

Once the bioimpedance value settled to a stable baseline, data was collected from the subject for ten minutes. The reconfigurable front end was programmed to alternate between the bipolar and tetrapolar configuration every minute, yielding five bipolar and five tetrapo-

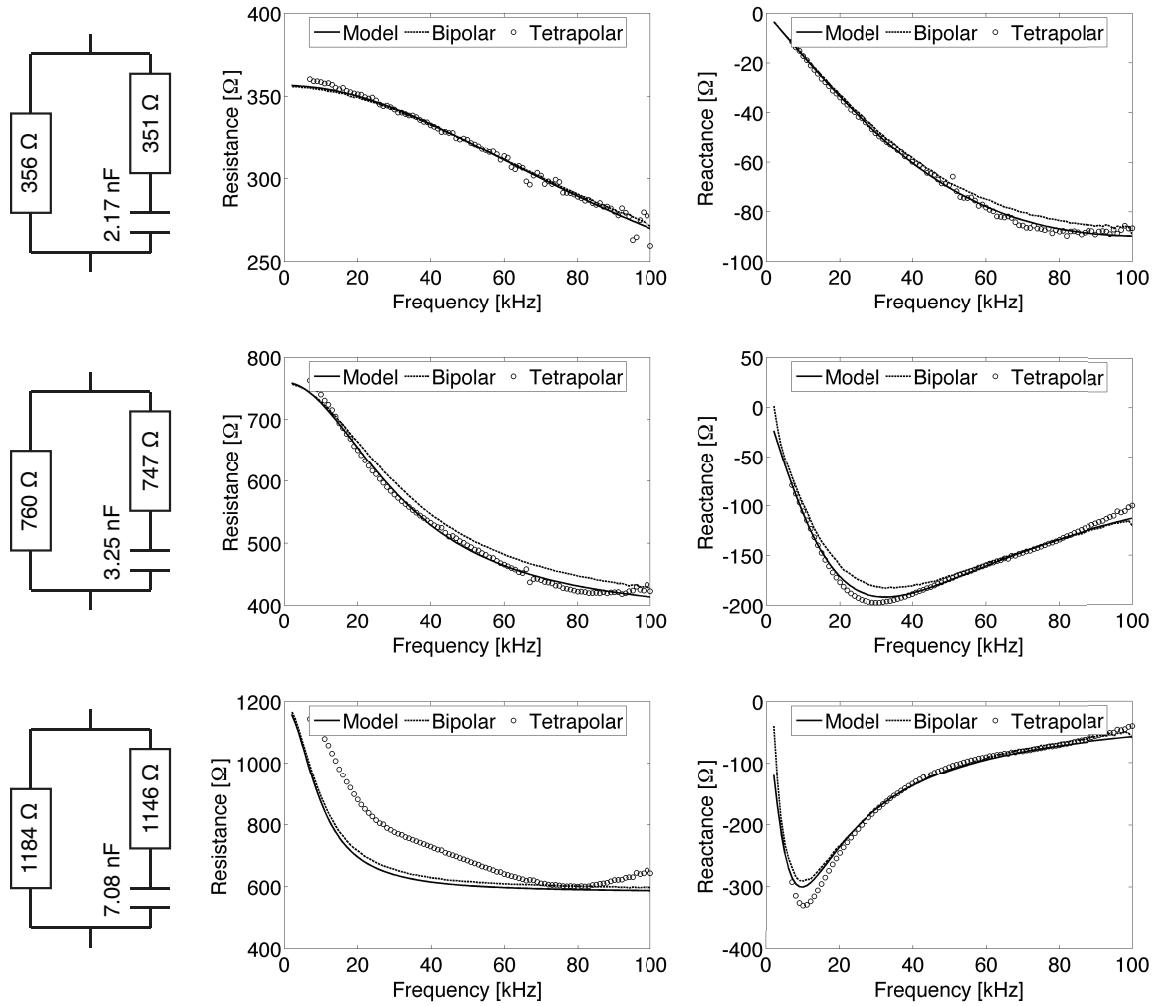


Figure 4.14: Measurement results of resistance and reactance from the wearable device configured for both bipolar (dashed line) and tetrapolar (marked with circles). Low, medium, and high impedance values are shown at the top, middle, and bottom, respectively.

lar sample runs per subject. A consistent offset (between approximately 100-500 Ω depending on the subject) was present in the wearable monitor during testing when the EBI100C was running. This offset was removed and data from the monitor and the EBI100C were passed through a simple moving average filter (with a moving window length of a fifth of second in both cases) to remove noise. These final filtered values were then compared.

4.2.3 The results

Results for low, medium, and high impedance values collected via the laboratory validation set up can be seen Figure 4.14. In all instances, the bipolar configuration (dashed line in the

Table 4.3: Measurement errors in resistance and reactance for each front end configuration without calibration, with a first order calibration, and with a second order calibration. Results are presented as an average with a 95% interval in brackets.

		Uncalibrated Error [%]	First Order Calibration Error [%]	Second Order Calibration Error [%]
Bipolar	R	-1.79 [-10.1, 6.47]	0.01 [-1.82 1.84]	0.006 [-1.04 1.06]
	X	1.76 [-21.8 25.3]	-0.07 [-12.2 12.1]	0.086 [-1.45 1.62]
Tetrapolar	R	5.74 [-18.7 30.2]	0.24 [-4.77 5.25]	0.11 [-3.47 3.69]
	X	6.39 [-0.38 13.2]	4.33 [-6.77 15.4]	3.1 [-1.56 7.76]

figure) gives very accurate results, often overlapping with the predicted value (solid line in the figure) from the model (seen to the left of the figure). Data from the tetrapolar configuration (marked with circles in the figure) show some random variation at low impedance values and a large offset in resistance at very high impedance values.

Table 4.3 demonstrates the precision of the reconfigurable monitor over the entire tested impedance range (minimum R = 193 Ω , maximum R = 1182 Ω , minimum X = -456 Ω , maximum X = -3.46 Ω). Even without calibration, the average error in the metric (defined here as the ratio of difference in measurement to the expected value) was between -1.79 and 1.76% for the bipolar configuration and between 5.74 and 6.39% for the tetrapolar configuration. Applying a first order calibration reduced the average error in resistance to less than 1% in both the tetrapolar and bipolar modes. While the average error in reactance was reduced in the bipolar case to -0.07%, the tetrapolar configuration produced an average error of 4.33%. The range of the errors was reduced in all instances for the first order calibration, with the exception of reactance as measured by the tetrapolar configuration. A second order calibration was also applied to the measurement, reducing the magnitude of the range of all errors to less than 10% over all measured impedances. The average error in resistance measurements was reduced to 0.006% and 0.11% in both the bipolar and tetrapolar mode, respectively, and the average error in reactance was reduced to 0.086% and 3.1% in the bipolar and tetrapolar mode, respectively. A typical example of uncalibrated (dashed line) and calibrated (solid line) resistance and reactance measurements from both the bipolar (blue) and the tetrapolar (red) configuration at the middle impedance presented in Figure 4.14 can be seen in Figure 4.15.

Figure 4.16 shows how the calibrated front end compares to a commercially available device by plotting the change in the magnitude of impedance over time for three subjects. In this figure, impedance is measured simultaneously by the reconfigurable front end and the EBI100C (solid line). For the first ten seconds impedance is measured by the reconfigurable




Figure 4.15: A typical resistance and reactance plot for the equivalent circuit measurements before (dashed lines) and after (solid lines) calibration for the bipolar (blue) and tetrapolar (red) configurations. The predicted response of the circuit (the one shown in the middle of Figure 4.14) is shown in black. In some areas it may be hard to see either the bipolar or tetrapolar response because of how well they align with predictions.

front end in the bipolar mode (blue line) and in the second ten seconds in the tetrapolar mode (red line). In all cases, the tracings show good agreement, with the wearable monitor results differing by less than $300 \text{ m}\Omega$ compared to the EBI100C over the course of the sample. The average error in the overall magnitude of impedance across all six subjects and all five individual trial runs was $1.56 \pm 5.42\%$ and $2.76 \pm 6.68\%$ in the bipolar and tetrapolar modes, respectively. The average error in the change in magnitude of impedance across all six subjects was comparable at $1.01 \pm 4.16\%$ and $2.71 \pm 4.53\%$. No significant difference in magnitude was observed between the bipolar and tetrapolar configurations, suggesting that both give equivalent results.

As mentioned in Chapter 3, Figure 4.16 also demonstrates a phenomenon while measuring subjects' continuous bioimpedance signal. In the top figure the signal varies largely in response to the respiratory cycle, whereas the bottom figure highlights a bioimpedance cycle almost exclusively dominated by the cardiac cycle. The figure in the middle shows a subject between these two ends of the spectrum, showing a superposition of both respiratory and cardiac signals. The clinical significance of such phenomena is currently unknown.

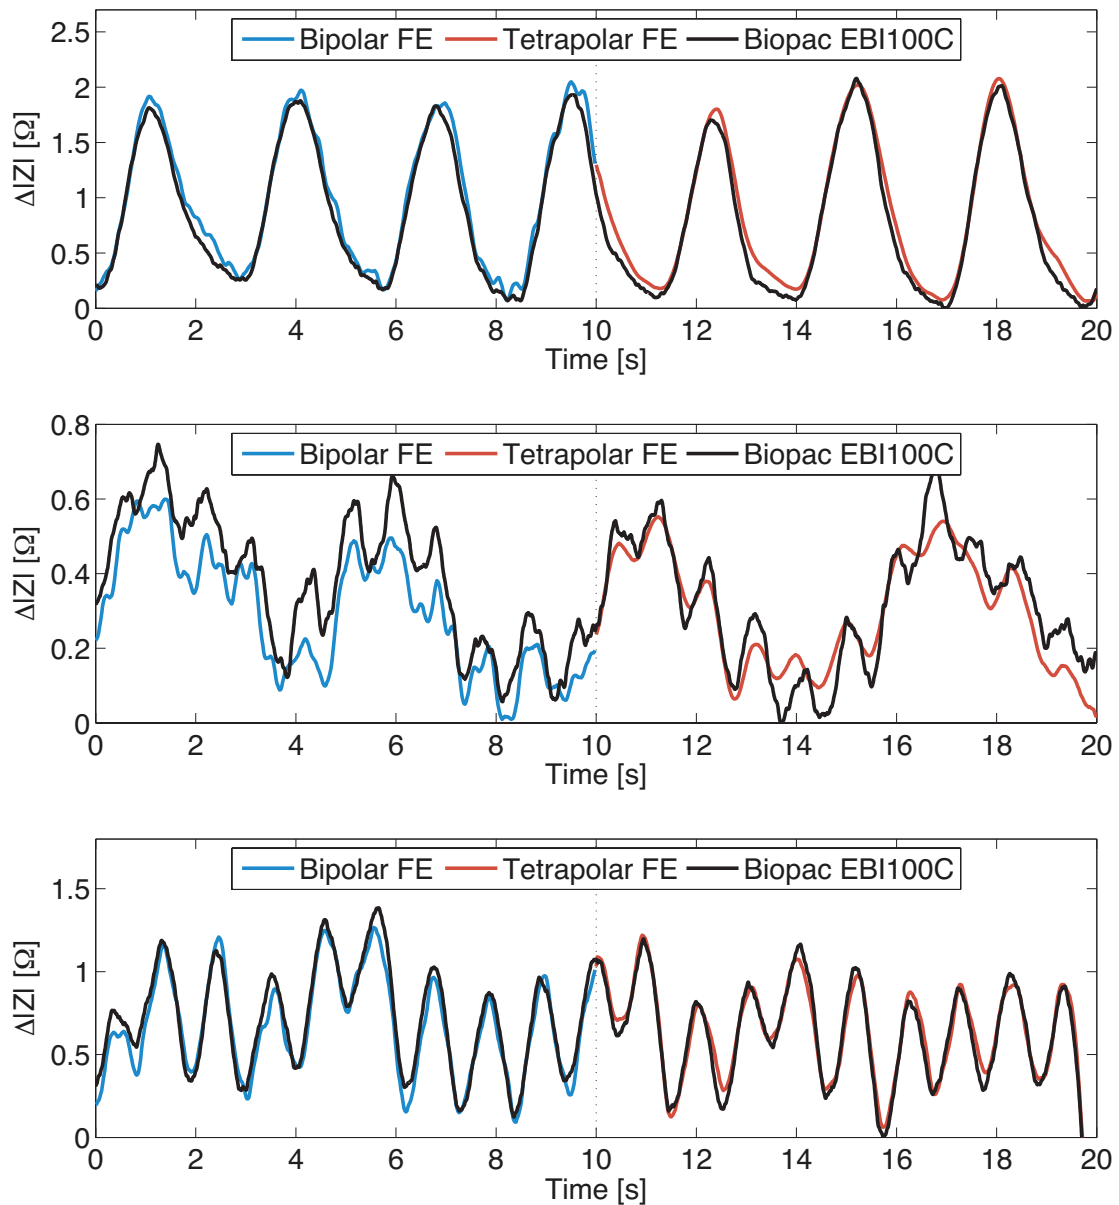


Figure 4.16: Results of the wearable device on three human subjects. At the 10 second mark of each graph the reconfigurable front end switched from the bipolar to the tetrapolar arrangement. (Moderate switch bouncing was observed and removed from this graph for the sake of clarity.) The top subject shows a response dominated by the respiratory signal with little evidence of a cardiac component. The bottom subject shows an opposite response, one where the cardiac signal dominates over respiration. In the middle of the two, both figuratively and literally, is a subject with both a cardiac and respiratory response. In all cases, the wearable device's results agreed well with those found via the EBI100C.

4.2.4 Conclusions

Wearable monitors capable of providing information relevant for clinical decision making will be an area of intense progress for years to come. Throughout this chapter I have outlined the design, calibration, and validation of a wearable bioimpedance measurement device for intravascular volume status monitoring. It is to my knowledge the first reported of its kind. While the measurement of fluid status and responsiveness remains hotly disputed, these measures lie at the foundation of treatments for many patients. That there exists a need to rapidly and accurately assess the volume status and responsiveness of a patient through many echelons of care is undeniable. That I have developed a monitor to serve such a role, I hope to have demonstrated.¹² How this helps solve the problem will only be seen with time.

¹²Others are vigorously working to develop their own wearable bioimpedance monitors. Several researchers continue to refine their own implementations [234] and even a few commercial devices are hitting the market are beginning to hit the market with these capabilities. For instance, AliphCom, that is “Jawbone”, has released devices (and patent applications [235]) wishing to use bioimpedance as a modality of import.

CHAPTER 5

The shape of things to come

We take our physiology with us as far as we can. Unchecked, unmeasured, untreated, it can take us in a moment—a lifetime too short. From the vigor of youth when health is in the soul to the long days of old age when its specter may overcast all else, knowing how our bodies work is vital to our very being. Determining if and when and how they fail allows us to treat them earlier and better and more optimally. We live in a time where more people are healthy than ever and we are more capable of treating the sickest among us. The spectrum of health of our populations ever increases with our knowledge of physiology and our ability to monitor and modify its function.

My research into two modalities for hemodynamic monitoring (bioimpedance and ultrasound) has shown me many places where we as a community of medical researchers and practitioners excel and stumble. Our collective strengths lie in innovation, in good research conducted by focused scientists, in the periodic cross-pollination of ideas from one field to the next. I would like to believe that in some small way this work may be a thread in such a tapestry. In fact, if I may be so brash, I see three distinct threads from the spools of this work.

The first and foremost contribution is to the science itself. With the help of many others, I was able to apply fundamental theories of dielectric materials to measure whole-body physiological events. Bridging the scientific gap between these two fields required a novel melding of mathematics, signal processing, and biology to translate the functions of the body to a new signal. A new vital sign. Moreover, once transduced—once translated from biochemistry to data—we as researchers and scientists can make predictions about its behavior in response to other factors. That is, we can explore entirely new vistas of physiology. I believe my work provides evidence that we can do so accurately and with confidence.

As previously alluded to, through the course of this work we observed many as of yet unreported phenomena including a variable cardiac signal found when measuring both

the inferior vena cava with ultrasound and in the limb with impedance plethysmography. The former is a venous response, one whose bedrock we have used as the foundation to test impedimetric intravascular volume evaluation. However, the extent to which the latter itself represent the venous system is an open debate, even among my collaborators and I. For my part, I believe the limb impedance reflects almost exclusively an arterial contribution to the signal, as is made quite clear when the derivative of the signal is taken, resulting in a waveform characteristically arterial. Beyond that, there must undoubtedly be an arterial signal somewhere in the limb impedance plethysmograph—blood does not merely flow through the veins. As such, the argument is one of the extent to which the signal is arterial and venous and not whether the signal is arterial or venous. Empirical arbitration is needed to settle this.

The second notable contribution of this work is the near future clinical implications. As much of the first chapter emphasizes, the current state of measuring intravascular volume within a modern hospital is woefully inadequate. With regularly used standards akin to flipping coins and new techniques coming to market requiring tea leaves to take readings, our patients deserve better, they need better. My work has demonstrated a technique that is better. I2VE's ability to correctly determine the level of IVC collapse is unprecedented (with large regions of 99% true positive accuracy) and a new wave in hydration status monitoring within the field of bioimpedance.

One crucial criticism here is the fact that no part of this work contains a clinical outcomes test or a measurement against another standard piece of equipment. Put another way, it does not mean much clinically if we can accurately measure a variable that has no relevance to the health of the patient. This is a fair critique and does represent a genuine weakness to the thesis I have put forward. But even as this work goes to print, those very experiments are being conducted and their results will one day be shared. If those results end up as optimistically as I predict they will, we will have solved the problem of measuring intravascular volume.¹ Until then, I must claim only that the results from this work can be tested more broadly by the community, knowing that the foundation has been set.

Confidence in a strong foundation arises when there is a minimum of surprises. That is, we must know when, where, and in what ways I2VE will work and when, where, and in what ways it will not. The assumptions necessary for its use and a few explicit limitations of its current implementation have already been laid out in Sections 3.2.2 and 3.3.5. But ultimately what is desired is an unambiguous demarcation of I2VE's abilities. Though

¹We would do well, while in this reverie, to remind ourselves that many researchers have thought precisely the same and have thus far met with disappointment. We should here at the end still maintain a humility that though we believe our efforts commendable they are only thus far deserving of a single, narrowly defined commendation.

the bounds of its total functionality have not yet been thoroughly explored, from what I have thus far discovered what follows is as clear and concise a border as I can draw. To work I2VE a patient must have a limb, the ability to breathe (either spontaneously or on a ventilator), a heart beat, a not-severely compromised cardiovascular circuit (hemorrhaging, for instance, would pose a large problem), and fluid within their vasculature. All of these details, as obvious as some are, are necessary for I2VE measurements to be made. If these conditions are not satisfied, the technique does not provide useful information. The most relevant case where I2VE may or may not work is in active and mobile patient populations, specifically those where one may wish to measure hydration status continuously in subjects that are highly active. Here, this distinction can be emphasized by considering subjects undergoing dialysis and those running marathons. Both experience large volume changes, yet the latter poses the harder metrological puzzle. For instance, when one lifts an arm in the air, the impedance plethysmographic signal will change through myriad factors including but not limited to electrode motion, electrode contact changes, and blood draining down the limb. In many cases, such changes dwarf the small signal we are attempting to observe and, would thus, currently fall within the categories of environments this technique is not appropriate for. However, with the rise of wearable electronics and the mad dash to correctly measure physiological signals on mobile populations, future researchers may yet be able to separate the confounding factors from the relevant ones and thus make I2VE viable more broadly. For now, I2VE is not appropriately suited to the domain of highly active subjects.

My third, final, and most hopeful contribution of this work is the opening of hemodynamic monitoring to a greater engineering approach. Hopefully the holistic conception and execution within this work shows that there are new approaches to be found in synthesizing the findings of different disciplines. There are places, though, where I faltered. Indeed, having tried to explain, measure, and work with the concept of intravascular volume status, I believe this is one area in need of a new paradigm. As we have seen, fluid therapy, one of the most widely used techniques in modern medicine, to this point has relied on techniques that are not too far from guessing in terms of accuracy and precision. Clinical decisions are unfortunately made with alarming frequency based on methods and modalities dependent on operator expertise, patient type, and sometimes even outdated physiological concepts.

One difficult question I have been posed often as I present this research to colleagues is “How would this data be presented in clinic?” That is, though I have a technique capable of measuring volume status and responsiveness with a high degree of accuracy and precision, how it might appear on a bedside monitor is not intuitive. An obvious solution might be to present a scrolling graph of impedance over time as is done for other measurements such

as electrocardiography and pulse oximetry. While I believe this may be worthwhile (it is, after all, an actual measurement of blood within the limbs), it would not show a patient's volume status. In fact, we ought to consider what we might mean when we want to display volume status. Currently, there is no unit of description. We can say when someone is over or under volume, and hazard a guess by what amount they are (usually as a function of "dry weight"), but with little precision. We can say if a person would or would not respond to an infusion of volume and even what type of volume expander to use to an extent,² but choices for amounts and rates are largely traditional or preferential.

What is needed is a unit. A unit of volume status and responsiveness that would say by what amount a patient needs (or does not need) intravascular volume. Such a unit may be tied to an ideal physiological end (a heart rate, a blood pressure), made patient specific (normalized to height, weight), and thought of in context (specific conditions or disease states may alter its scaling or offset). Though to an extent this is what current practice may aim to do, without the quantifiable the qualitative reigns. The hunch is followed when it is the only flickering light in the darkness. The gut leads the head. While there is much to be said for the intuition of the physician (their years of experience are akin to personal machine learning), there is as much to commend about precise, measurable, and predictable units. This may just be one engineer's take on the state of modern medicine, but I feel that unless and until we are serious about what we measure, we will never be able to do it well.

My own work must be improved. Many aspects of ultrasound ought to become automated so that talented physicians are not tracking pixels and all clinicians can pick up and use a device. Machines ought to recognize what they are imaging (an inferior vena cava or an arteriovenous fistula, for instance) and provide doctors with useful information on the fly: the diameters and distensibilities of vessels, anechoic regions of blood (possibly where it is not supposed to be), volume responsiveness, etc. This must be done so that we utilize our time and our signals to the best of our abilities.

Bioimpedance itself as a signal needs to address its susceptibility to motion artifact to be viable as a long term real world solution for measuring biological signals. Unlike other signals whose occurrence is not affected by a body in motion, the very nature of measuring blood within the body via limb impedance is fraught with artifacts both as a condition of physiology and measurement. Separating electrode motion from fluid motion is a difficult and not at all solved problem in the field. The current solution (one I myself have employed) is to either not measure during times of motion or to dampen all responses blindly. Neither is sufficient if we are to truly measure continuously.

All of this to say the main thrust of this work has been to measure more physiology

²Though, as with much in this field, this remains an area of on-going debate.

in more places. This is in line with the historical progression of health monitoring. From the millennia that passed where our species's sickness, disease, and injury were merely to be waited out, we have become a civilization where we view our clinical needs as on-demand. Annual medical exams, widespread around the clock urgent and emergent care, and the draw of several major life events (birth, suffering, death) have made health as a daily concern a reality. The time is not far off when the capability to measure every vital sign (every heart beat, every breath) from the moment of birth to the time of death will be possible. Fortunately, at the gestation of my techniques, their roles and limitations in the arena of constant monitoring have been recognized and designed into it. Of course, only future work will determine whether the tree I have planted here will bear fruit or merely serve as kindling in the pyres of scientific progress.

I believe the shape of things to come is multidimensional, multidisciplinary, and multimodal. We must not just consider a person in time and space and only at the bedside. They have a history, an anatomy, a place in this world all their own, and it will take teams of engineers, scientists, and physicians to figure out how best to understand it all to the benefit of that person. We must leverage all of the signals we measure—from ultrasound to bioimpedance, from electrocardiography to photoplethysmography—to detect as much as we can to help our patients. The physiology is there. We need only to listen carefully.

Bibliography

BIBLIOGRAPHY

- [1] “Global, regional, and national age-sex specific all-cause and cause-specific mortality for 240 causes of death , 1990-2013 : a systematic analysis for the Global Burden of Disease Study 2013,” *Lancet*, vol. 385, no. 9963, pp. 117–171, 2015.
- [2] Edwards Lifesciences LLC, “Normal Hemodynamic Parameters and Laboratory,” 2009.
- [3] D. J. Funk, E. Jacobsohn, and A. Kumar, “The role of venous return in critical illness and shock-part I: physiology.,” *Critical care medicine*, vol. 41, pp. 255–62, 2013.
- [4] M. Lichtwarck-Aschoff, R. Beale, and U. J. Pfeiffer, “Central venous pressure, pulmonary artery occlusion pressure, intrathoracic blood volume, and right ventricular end-diastolic volume as indicators of cardiac preload,” *Journal of Critical Care*, vol. 11, no. 4, pp. 180–188, 1996.
- [5] D. D. Glower, J. a. Spratt, N. D. Snow, J. S. Kabas, J. W. Davis, C. O. Olsen, G. S. Tyson, D. C. Sabiston, and J. S. Rankin, “Linearity of the Frank-Starling relationship in the intact heart: the concept of preload recruitable stroke work.,” *Circulation*, vol. 71, no. 5, pp. 994–1009, 1985.
- [6] D. T. Mangano, D. C. Van Dyke, and R. J. Ellis, “The effect of increasing preload on ventricular output and ejection in man. Limitations of the Frank-Starling Mechanism.,” *Circulation*, vol. 62, no. 3, pp. 535–541, 1980.
- [7] R. H. Schwinger, M. Böhm, a. Koch, U. Schmidt, I. Morano, H. J. Eissner, P. Uberfuhr, B. Reichart, and E. Erdmann, “The failing human heart is unable to use the Frank-Starling mechanism.,” *Circulation research*, vol. 74, no. 5, pp. 959–969, 1994.
- [8] N. Rott, “Note on the history of the the Reynolds Number,” *Annual Review of Fluid Mechanics*, vol. 22, pp. 1–11, 1990.

- [9] G. G. Stokes, "On the Effect of the Internal Friction of Fluids on the Motion of Pendulums," *Transactions of the Cambridge Philosophical Society*, vol. 9, 1851.
- [10] D. N. Ku, "Blood flow in arteries," pp. 399–434, 1997.
- [11] J. R. Womersley, "Method for the calculation of velocity, rate of flow and viscous drag in arteries when the pressure gradient is known.," *The Journal of physiology*, vol. 127, no. 3, pp. 553–563, 1955.
- [12] P. Perera, T. Mailhot, D. Riley, and D. Mandavia, "The RUSH Exam : Rapid Ultrasound in SHock in the Evaluation of the Critically Ill," *Shock*, vol. 28, pp. 29–56, 2010.
- [13] P. Perera, T. Mailhot, D. Riley, and D. Mandavia, "The RUSH exam 2012: Rapid ultrasound in shock in the evaluation of the critically ill patient," *Ultrasound Clinics*, vol. 7, no. 2, pp. 255–278, 2012.
- [14] M. Cecconi, A. K. Parsons, and A. Rhodes, "What is a fluid challenge?," *Current opinion in critical care*, vol. 17, no. 3, pp. 290–295, 2011.
- [15] M. Antonelli, M. Levy, P. J. D. Andrews, J. Chastre, L. D. Hudson, C. Manthous, G. U. Meduri, R. P. Moreno, C. Putensen, T. Stewart, and A. Torres, "Hemodynamic monitoring in shock and implications for management. International Consensus Conference, Paris, France, 27-28 April 2006.," *Intensive care medicine*, vol. 33, no. 4, pp. 575–590, 2007.
- [16] J. G. Jones and C. a. Wardrop, "Measurement of blood volume in surgical and intensive care practice.," *British journal of anaesthesia*, vol. 84, no. 2, pp. 226–235, 2000.
- [17] J. V. Nixon, R. G. Murray, P. D. Leonard, J. H. Mitchell, and C. G. Blomqvist, "Effect of large variations in preload on left ventricular performance characteristics in normal subjects.," *Circulation*, vol. 65, no. 4, pp. 698–703, 1982.
- [18] P. E. Marik, M. Baram, and B. Vahid, "Does central venous pressure predict fluid responsiveness?," *Chest*, vol. 134, no. 1, pp. 172–178, 2008.
- [19] P. E. Marik, "Techniques for assessment of intravascular volume in critically ill patients.," *Journal of intensive care medicine*, vol. 24, pp. 329–337, 2009.
- [20] W. Frank Peacock and K. M. Soto, "Current technique of fluid status assessment," *Congestive Heart Failure*, vol. 16, no. SUPPL. 1, pp. 45–51, 2010.

- [21] K. Kalantari, J. N. Chang, C. Ronco, and M. H. Rosner, "Assessment of intravascular volume status and volume responsiveness in critically ill patients.," *Kidney international*, vol. 83, no. 6, pp. 1017–28, 2013.
- [22] S. McGee, "Is This Patient Hypovolemic?," *JAMA: The Journal of the American Medical Association*, vol. 281, pp. 1022–1029, 1999.
- [23] C. S. Wang, J. M. FitzGerald, M. Schulzer, E. Mak, and N. T. Ayas, "Does this dyspneic patient in the emergency department have congestive heart failure?," *JAMA : the journal of the American Medical Association*, vol. 294, no. 15, pp. 1944–1956, 2005.
- [24] P. R. Eisenberg, a. S. Jaffe, and D. P. Schuster, "Clinical evaluation compared to pulmonary artery catheterization in the hemodynamic assessment of critically ill patients.," 1984.
- [25] L. W. Stevenson and J. K. Perloff, "The limited reliability of physical signs for estimating hemodynamics in chronic heart failure.," *JAMA : the journal of the American Medical Association*, vol. 261, pp. 884–888, 1989.
- [26] E. P. Rivers, B. Nguyen, S. Havstad, J. Ressler, A. Muzzin, B. Knoblich, E. Peterson, and M. Tomlanovich, "Early goal-directed therapy in the treatment of severe sepsis and septic shock," *New England Journal of Medicine*, vol. 345, no. 19, pp. 1368–1377, 2001.
- [27] R. Dellinger, M. Levy, and A. Rhodes, "Surviving Sepsis Campaign: international guidelines for management of severe sepsis and septic shock, 2012," *Intensive Care Medicine*, vol. 41, no. 2, pp. 580–637, 2013.
- [28] S. Magder, "How to use central venous pressure measurements.," *Current opinion in critical care*, vol. 11, no. 3, pp. 264–270, 2005.
- [29] A. C. Guyton, "Determination of cardiac output by equating venous return curves with cardiac response curves.," *Physiological Reviews*, vol. 35, no. 1, pp. 123–9, 1955.
- [30] M. Cecconi, H. D. Aya, M. Geisen, C. Ebn, N. Fletcher, R. M. Grounds, and A. Rhodes, "Changes in the mean systemic filling pressure during a fluid challenge in postsurgical intensive care patients," *Intensive Care Medicine*, vol. 39, no. 7, pp. 1299–1305, 2013.

- [31] A. Barbeito and J. B. Mark, “Arterial and Central Venous Pressure Monitoring,” *Anesthesiology Clinics of North America*, vol. 24, no. 4, pp. 717–735, 2006.
- [32] C. M. Boltwood, “Ventricular performance related to transmural filling pressure in clinical cardiac tamponade.,” *Circulation*, vol. 75, no. 5, pp. 941–955, 1987.
- [33] K. K. Figg and E. C. Nemergut, “Error in central venous pressure measurement,” *Anesthesia and Analgesia*, vol. 108, no. 4, pp. 1209–1211, 2009.
- [34] P. E. Marik, M. Baram, and B. Vahid, “Does central venous pressure predict fluid responsiveness?: A Systematic Review of the Literature and the Tale of Seven Mares,” *Chest*, vol. 134, pp. 172–178, 2008.
- [35] P. E. Marik and R. Cavallazzi, “Does the central venous pressure predict fluid responsiveness? An updated meta-analysis and a plea for some common sense.,” *Critical care medicine*, vol. 41, no. 7, pp. 1774–81, 2013.
- [36] A. Kumar, R. Anel, E. Bunnell, K. Habet, S. Zanotti, S. Marshall, A. Neumann, A. Ali, M. Cheang, C. Kavinsky, and J. E. Parrillo, “Pulmonary artery occlusion pressure and central venous pressure fail to predict ventricular filling volume, cardiac performance, or the response to volume infusion in normal subjects.,” *Critical care medicine*, vol. 32, no. 3, pp. 691–699, 2004.
- [37] S. Gelman, “Venous Function and Central Venous Pressure,” *Anesthesiology*, vol. 108, no. 4, pp. 735–48, 2008.
- [38] F. Michard and J.-L. Teboul, “Predicting fluid responsiveness in ICU patients: a critical analysis of the evidence.,” *Chest*, vol. 121, pp. 2000–2008, 2002.
- [39] D. Osman, C. Ridel, P. Ray, X. Monnet, N. Anguel, C. Richard, and J.-L. Teboul, “Cardiac filling pressures are not appropriate to predict hemodynamic response to volume challenge.,” *Critical care medicine*, vol. 35, no. 1, pp. 64–68, 2007.
- [40] W. H. Gaasch, W. E. Battle, A. a. Oboler, J. S. Banas, and H. J. Levine, “Left Ventricular Stress and Compliance in Man,” vol. XLV, no. April 1972, 1972.
- [41] W. H. Gaasch, M. a. Quinones, E. Waisser, H. G. Thiel, and J. K. Alexander, “Diastolic compliance of the left ventricle in man.,” *The American journal of cardiology*, vol. 36, no. 2, pp. 193–201, 1975.
- [42] B. S. Lewis and M. S. Gotsman, “Current concepts of left ventricular relaxation and compliance.,” *American heart journal*, vol. 99, no. 1, pp. 101–112, 1980.

- [43] J. A. Martyn, M. T. Snider, L. F. Farago, and J. F. Burke, "Thermodilution right ventricular volume: a novel and better predictor of volume replacement in acute thermal injury.," *The Journal of trauma*, vol. 21, no. 8, pp. 619–626, 1981.
- [44] C. Reuse, J. L. Vincent, and M. R. Pinsky, "Measurements of right ventricular volumes during fluid challenge," *Chest*, vol. 98, no. 6, pp. 1450–1454, 1990.
- [45] L. N. Diebel, R. F. Wilson, M. G. Tagett, and R. a. Kline, "End-diastolic volume. A better indicator of preload in the critically ill.," *Archives of surgery*, vol. 127, no. 7, pp. 817–821; discussion 821–822, 1992.
- [46] J. G. Wagner and J. W. Leatherman, "Right ventricular end-diastolic volume as a predictor of the hemodynamic response to a fluid challenge," *Chest*, vol. 113, no. 4, pp. 1048–1054, 1998.
- [47] K. Scheuren, M. N. Wentz, C. Hainer, M. Scheffler, C. Lichtenstern, E. Martin, J. Schmidt, C. Bopp, and M. a. Weigand, "Left ventricular end-diastolic area is a measure of cardiac preload in patients with early septic shock.," *European journal of anaesthesiology*, vol. 26, no. 9, pp. 759–765, 2009.
- [48] D. M. Thys, Z. Hillel, M. E. Goldman, B. P. Mindich, and J. A. Kaplan, "A comparison of hemodynamic indices derived by invasive monitoring and two-dimensional echocardiography.," *Anesthesiology*, vol. 67, no. 5, pp. 630–634, 1987.
- [49] D. a. Reuter, T. W. Felbinger, C. Schmidt, E. Kilger, O. Goedje, P. Lamm, and A. E. Goetz, "Stroke volume variations for assessment of cardiac responsiveness to volume loading in mechanically ventilated patients after cardiac surgery," *Intensive Care Medicine*, vol. 28, no. 4, pp. 392–398, 2002.
- [50] M. Feissel, F. Michard, I. Mangin, O. Ruyter, and J.-p. Faller, "Respiratory Changes in Aortic Blood Velocity as an Indicator of Fluid Responsiveness in Ventilated Patients With Septic Shock," *Chest*, vol. 119, no. 3, pp. 867–873, 2001.
- [51] L. Belloni, A. Pisano, A. Natale, M. R. Piccirillo, L. Piazza, G. Ismeno, and G. De Martino, "Assessment of Fluid-Responsiveness Parameters for Off-Pump Coronary Artery Bypass Surgery: A Comparison Among LiDCO, Transesophageal Echocardiography, and Pulmonary Artery Catheter," *Journal of Cardiothoracic and Vascular Anesthesia*, vol. 22, no. 2, pp. 243–248, 2008.

- [52] S. a. Josephs, “The use of current hemodynamic monitors and echocardiography in resuscitation of the critically ill or injured patient.,” *International anesthesiology clinics*, vol. 45, no. 3, pp. 31–59, 2007.
- [53] G. S. Mintz, M. N. Kotler, W. R. Parry, a. S. Iskandrian, and S. a. Kane, “Real-time inferior vena caval ultrasonography: normal and abnormal findings and its use in assessing right-heart function.,” *Circulation*, vol. 64, pp. 1018–1025, 1981.
- [54] J. Jue, W. Chung, and N. B. Schiller, “Does Inferior Vena Cava Size Predict Right Atrial Pressures in Patients Receiving Mechanical Ventilation?,” *Journal of the American Society of Echocardiography*, vol. 5, pp. 613–619, nov 1992.
- [55] Y. Yanagawa, K. Nishi, T. Sakamoto, and Y. Okada, “Early diagnosis of hypovolemic shock by sonographic measurement of inferior vena cava in trauma patients.,” *The Journal of trauma*, vol. 58, no. 4, pp. 825–829, 2005.
- [56] Y. Yanagawa, T. Sakamoto, and Y. Okada, “Hypovolemic shock evaluated by sonographic measurement of the inferior vena cava during resuscitation in trauma patients.,” *The Journal of trauma*, vol. 63, no. 6, pp. 1245–1248; discussion 1248, 2007.
- [57] W. F. Peacock, “Use of bioimpedance vector analysis in critically ill and cardiorenal patients,” *Contributions to Nephrology*, vol. 165, pp. 226–235, 2010.
- [58] S. D. Somma, H. C. Lukaski, and M. Codognotto, “Consensus paper on the use of BIVA (Bioelectrical Impedance Vector Analysis) in medicine for the management of body hydration,” no. Figure 1, 2011.
- [59] W. W. Tang, V. K. Chopra, N. Chakravarthy, I. Libbus, and R. P. Katra, “External Wireless Monitoring of Bioimpedance in Heart Failure Patients: Results from the ACUTE Study,” *Journal of Cardiac Failure*, vol. 16, pp. S70–S71, aug 2010.
- [60] P. E. Marik, R. Cavallazzi, T. Vasu, and A. Hirani, “Dynamic changes in arterial waveform derived variables and fluid responsiveness in mechanically ventilated patients: a systematic review of the literature.,” *Critical care medicine*, vol. 37, no. 9, pp. 2642–2647, 2009.
- [61] F. Michard and J. L. Teboul, “Using heart-lung interactions to assess fluid responsiveness during mechanical ventilation.,” *Critical Care*, vol. 4, no. 5, pp. 282–289, 2000.

- [62] L. J. Montenij, E. E. de Waal, and W. F. Buhre, "Arterial waveform analysis in anesthesia and critical care," *Current Opinion in Anaesthesiology*, vol. 24, no. 6, pp. 551–556, 2011.
- [63] M. Cannesson, Y. Attof, P. Rosamel, O. Desebbe, P. Joseph, O. Metton, O. Bastien, and J.-J. Lehot, "Respiratory variations in pulse oximetry plethysmographic waveform amplitude to predict fluid responsiveness in the operating room.," *Anesthesiology*, vol. 106, no. 6, pp. 1105–1111, 2007.
- [64] O. Desebbe and M. Cannesson, "Using ventilation-induced plethysmographic variations to optimize patient fluid status.," *Current opinion in anaesthesiology*, vol. 21, no. 6, pp. 772–778, 2008.
- [65] a. Perner and T. Faber, "Stroke volume variation does not predict fluid responsiveness in patients with septic shock on pressure support ventilation," *Acta Anaesthesiologica Scandinavica*, vol. 50, pp. 1068–1073, 2006.
- [66] S. Soubrier, F. Saulnier, H. Hubert, P. Delour, H. Lenci, T. Onimus, S. Nseir, and A. Durocher, "Can dynamic indicators help the prediction of fluid responsiveness in spontaneously breathing critically ill patients?," *Intensive Care Medicine*, vol. 33, no. 7, pp. 1117–1124, 2007.
- [67] P. M. Dark and M. Singer, "The validity of trans-esophageal Doppler ultrasonography as a measure of cardiac output in critically ill adults," *Intensive Care Medicine*, vol. 30, no. 11, pp. 2060–2066, 2004.
- [68] C. Barbier, Y. Loubières, C. Schmit, J. Hayon, J. L. Ricôme, F. Jardin, and A. Vieillard-Baron, "Respiratory changes in inferior vena cava diameter are helpful in predicting fluid responsiveness in ventilated septic patients," *Intensive Care Medicine*, vol. 30, no. 9, pp. 1740–1746, 2004.
- [69] A. Dipti, Z. Soucy, A. Surana, and S. Chandra, "Role of inferior vena cava diameter in assessment of volume status: A meta-analysis," *American Journal of Emergency Medicine*, vol. 30, no. 8, 2012.
- [70] M. Feissel, F. Michard, J. P. Faller, and J. L. Teboul, "The respiratory variation in inferior vena cava diameter as a guide to fluid therapy," *Intensive Care Medicine*, vol. 30, pp. 1834–1837, 2004.
- [71] P. Ferrada, R. J. Anand, J. Whelan, M. a. Aboutanos, T. Duane, A. Malhotra, and R. Ivatury, "Qualitative assessment of the inferior vena cava: useful tool for the

evaluation of fluid status in critically ill patients.,” *The American surgeon*, vol. 78, pp. 468–70, 2012.

- [72] L. Muller, X. Bobbia, M. Toumi, G. Louart, N. Molinari, B. Ragonnet, H. Quintard, M. Leone, L. Zoric, and J. Y. Lefrant, “Respiratory variations of inferior vena cava diameter to predict fluid responsiveness in spontaneously breathing patients with acute circulatory failure: need for a cautious use.,” *Critical care (London, England)*, vol. 16, p. R188, jan 2012.
- [73] J. Jabot, J. L. Teboul, C. Richard, and X. Monnet, “Passive leg raising for predicting fluid responsiveness: Importance of the postural change,” *Intensive Care Medicine*, vol. 35, no. 1, pp. 85–90, 2009.
- [74] a. Lafanechère, F. Pène, C. Goulenok, a. Delahaye, V. Mallet, G. Choukroun, J. D. Chiche, J. P. Mira, and a. Cariou, “Changes in aortic blood flow induced by passive leg raising predict fluid responsiveness in critically ill patients.,” *Critical care (London, England)*, vol. 10, no. 5, p. R132, 2006.
- [75] J.-L. Teboul and X. Monnet, “Prediction of volume responsiveness in critically ill patients with spontaneous breathing activity.,” *Current opinion in critical care*, vol. 14, no. 3, pp. 334–339, 2008.
- [76] X. Monnet, M. Rienzo, D. Osman, N. Anguel, C. Richard, M. R. Pinsky, and J.-L. Teboul, “Passive leg raising predicts fluid responsiveness in the critically ill.,” *Critical care medicine*, vol. 34, no. 5, pp. 1402–1407, 2006.
- [77] X. Monnet, C. Richard, and J. L. Teboul, “Passive leg raising,” *Intensive Care Medicine: Annual Update 2007*, pp. 542–548, 2007.
- [78] F. Cavallaro, C. Sandroni, C. Marano, G. La Torre, A. Mannocci, C. De Waure, G. Bello, R. Maviglia, and M. Antonelli, “Diagnostic accuracy of passive leg raising for prediction of fluid responsiveness in adults: Systematic review and meta-analysis of clinical studies,” *Intensive Care Medicine*, vol. 36, no. 9, pp. 1475–1483, 2010.
- [79] K. Lakhal, S. Ehrmann, I. Runge, D. Benzekri-Lefèvre, A. Legras, P. F. Dequin, E. Mercier, M. Wolff, B. Régnier, and T. Boulain, “Central venous pressure measurements improve the accuracy of leg raising-induced change in pulse pressure to predict fluid responsiveness,” *Intensive Care Medicine*, vol. 36, no. 6, pp. 940–948, 2010.

- [80] N. Y. Raval, P. Squara, M. Cleman, K. Yalamanchili, M. Winklmaier, and D. Burkhoff, "Multicenter evaluation of noninvasive cardiac output measurement by bioreactance technique," *Journal of Clinical Monitoring and Computing*, vol. 22, no. 2, pp. 113–119, 2008.
- [81] D. H. Conway, O. A. Hussain, and I. Gall, "A comparison of noninvasive bioreactance with oesophageal Doppler estimation of stroke volume during open abdominal surgery," *European Journal of Anaesthesiology*, vol. 30, no. 8, pp. 501–508, 2013.
- [82] D. Fagnoul, J.-L. Vincent, and D. D. Backer, "Cardiac output measurements using the bioreactance technique in critically ill patients.," *Critical care (London, England)*, vol. 16, no. 6, p. 460, 2012.
- [83] B. Benomar, A. Ouattara, P. Estagnasie, A. Brusset, and P. Squara, "Fluid responsiveness predicted by noninvasive bioreactance-based passive leg raise test," *Intensive Care Medicine*, vol. 36, no. 11, pp. 1875–1881, 2010.
- [84] S. Marqué, A. Cariou, J.-D. Chiche, and P. Squara, "Comparison between FloTrac-Vigileo and Bioreactance, a totally noninvasive method for cardiac output monitoring.," *Critical care (London, England)*, vol. 13, no. 3, p. R73, 2009.
- [85] P. E. Marik, A. Levitov, A. Young, and L. Andrews, "The use of bioreactance and carotid doppler to determine volume responsiveness and blood flow redistribution following passive leg raising in hemodynamically unstable patients," *Chest*, vol. 143, no. 2, pp. 364–370, 2013.
- [86] E. Kupersztych-Hagege, J. L. Teboul, a. Artigas, a. Talbot, C. Sabatier, C. Richard, and X. Monnet, "Bioreactance is not reliable for estimating cardiac output and the effects of passive leg raising in critically ill patients," *British Journal of Anaesthesia*, vol. 111, no. 6, pp. 961–966, 2013.
- [87] J. M. Fitzpatrick, D. L. G. Hill, and C. R. Maurer, "Image Registration," *Handbook of Medical Imaging, Volume2: Medical Image Processing and Analysis*, pp. 447–514, 2000.
- [88] B. Zitová and J. Flusser, "Image registration methods: A survey," *Image and Vision Computing*, vol. 21, no. 11, pp. 977–1000, 2003.
- [89] J. B. Maintz and M. A. Viergever, "A survey of medical image registration.," *Medical image analysis*, vol. 2, no. 1, pp. 1–36, 1998.

- [90] J. Li, N. M. Allinson, J. L. A., and N. M. Allinson, “A comprehensive review of current local features for computer vision,” *Neurocomputing*, vol. 71, no. 10-12, pp. 1771–1787, 2008.
- [91] G. Stockman, S. Kopstein, and S. Benett, “Matching Images to Models for Registration and Object Detection via Clustering.,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 4, no. 3, pp. 229–241, 1982.
- [92] C. Harris and M. Stephens, “A Combined Corner and Edge Detector,” *Proceedings of the Alvey Vision Conference 1988*, pp. 23.1–23.6, 1988.
- [93] S. Banerjee, D. Mukherjee, and D. Majumdar, “Point landmarks for registration of CT and MR images,” *Pattern Recognition Letters*, vol. 16, no. 10, pp. 1033–1042, 1995.
- [94] S. Bhardwaj and A. Mittal, “A Survey on Various Edge Detector Techniques,” *Procedia Technology*, vol. 4, pp. 220–226, 2012.
- [95] L. S. Davis, “A survey of edge detection techniques,” *Computer Graphics and Image Processing*, vol. 4, no. 3, pp. 248–270, 1975.
- [96] N. R. Pal and S. K. Pal, “A review on image segmentation techniques,” *Pattern Recognition*, vol. 26, no. 9, pp. 1277–1294, 1993.
- [97] K. Fu and J. Mui, “A survey on image segmentation,” *Pattern Recognition*, vol. 13, no. 1, pp. 3–16, 1981.
- [98] T. S. Caetano, T. Caelli, D. Schuurmans, and D. a. C. Barone, “Graphical models and point pattern matching.,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 28, no. 10, pp. 1646–63, 2006.
- [99] A. Goshtasby and G. C. Stockman, “Point pattern matching using convex hull edges,” *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-15, no. 5, pp. 631–637, 1985.
- [100] A. A. Goshtasby, *Image Registration: Principles, Tools and Methods*. 2012.
- [101] L. Zagorchev and A. Goshtasby, “A comparative study of transformation functions for nonrigid image registration.,” *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society*, vol. 15, no. 3, pp. 529–538, 2006.

- [102] A. Goshtasby, "Image registration by local approximation methods," *Image and Vision Computing*, vol. 6, no. 4, pp. 255–261, 1988.
- [103] A. Goshtasby, "Piecewise linear mapping functions for image registration," *Pattern Recognition*, vol. 19, no. 6, pp. 459–466, 1986.
- [104] A. Goshtasby, "Piecewise cubic mapping functions for image registration," *Pattern Recognition*, vol. 20, no. 5, p. 525, 1987.
- [105] R. Bajcsy and S. Kovačič, "Multiresolution elastic matching," *Computer Vision, Graphics, and Image Processing*, vol. 46, no. 1, pp. 1–21, 1989.
- [106] M. Bro-Nielsen and C. Gramkow, "Fast Fluid Registration of medical images," vol. 25, pp. 265–276, 1996.
- [107] J. L. Barron, D. J. Fleet, and S. S. Beauchemin, "Performance of optical flow techniques," *International Journal of Computer Vision*, vol. 12, no. 1, pp. 43–77, 1994.
- [108] A. Singh, "An Estimation-Theoretic Framework for Image-Flow Computation," *Proceedings, Third International Conference on Computer Vision, 1990.*, 1990.
- [109] A. Singh and P. Allen, "Image-Flow Computation: An Estimation-Theoretic and a Unified Perspective," *CVGIP: Image Understanding*, vol. 56, no. 2, pp. 152–177, 1992.
- [110] B. K. Horn and B. G. Schunck, "Determining optical flow," *Artificial Intelligence*, vol. 17, no. 1-3, pp. 185–203, 1981.
- [111] B. D. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision," 1981.
- [112] A. Barjatya, "Block Matching Algorithms For Motion Estimation," *IEEE Transactions Evolution Computation*, vol. 8, no. 3, pp. 225–229, 2004.
- [113] J. S. J. Shi and C. Tomasi, "Good features to track," *Computer Vision and Pattern Recognition, 1994. Proceedings CVPR '94., 1994 IEEE Computer Society Conference on*, pp. 593–600, 1994.
- [114] C. Harris and M. Stephens, "A Combined Corner and Edge Detector," *Proceedings of the Alvey Vision Conference 1988*, pp. 147–151, 1988.

- [115] “KDOQI Clinical Practice Guidelines and Clinical Practice Recommendations for Diabetes and Chronic Kidney Disease,” *American Journal of Kidney Diseases*, vol. 49, no. 2 SUPPL. 2, 2007.
- [116] H. C. Rayner, A. Besarab, W. W. Brown, A. Disney, A. Saito, and R. L. Pisoni, “Vascular access results from the Dialysis Outcomes and Practice Patterns Study (DOPPS): Performance against Kidney Disease Outcomes Quality Initiative (K/DOQI) Clinical Practice Guidelines,” *American Journal of Kidney Diseases*, vol. 44, pp. 22–26, nov 2004.
- [117] J. Ethier, D. C. Mendelssohn, S. J. Elder, T. Hasegawa, T. Akizawa, T. Akiba, B. J. Canaud, and R. L. Pisoni, “Vascular access use and outcomes: An international perspective from the dialysis outcomes and practice patterns study,” *Nephrology Dialysis Transplantation*, vol. 23, no. 10, pp. 3219–3226, 2008.
- [118] S. T. Patel, J. Hughes, J. L. Mills, and T. S. Huber, “Failure of arteriovenous fistula maturation: An unintended consequence of exceeding Dialysis Outcome Quality Initiative guidelines for hemodialysis access,” *Journal of Vascular Surgery*, vol. 38, no. 3, pp. 439–445, 2003.
- [119] K. D. Gibson, D. L. Gillen, M. T. Caps, T. R. Kohler, D. J. Sherrard, and C. O. Stehman-Breen, “Vascular access survival and incidence of revisions: A comparison of prosthetic grafts, simple autogenous fistulas, and venous transposition fistulas from the United States Renal Data System Dialysis Morbidity and Mortality Study,” *Journal of Vascular Surgery*, vol. 34, no. 4, pp. 694–700, 2001.
- [120] C. E. Lok, J. M. Sontrop, G. Tomlinson, D. Rajan, M. Cattral, G. Oreopoulos, J. Harris, and L. Moist, “Cumulative patency of contemporary fistulas versus grafts (2000-2010),” *Clinical Journal of the American Society of Nephrology*, vol. 8, no. 5, pp. 810–818, 2013.
- [121] J. Parmar, M. Aslam, and N. Standfield, “Pre-operative Radial Arterial Diameter Predicts Early Failure of Arteriovenous Fistula (AVF) for Haemodialysis,” *European Journal of Vascular and Endovascular Surgery*, vol. 33, no. 1, pp. 113–115, 2007.
- [122] R. R. Mendes, M. a. Farber, W. a. Marston, L. C. Dinwiddie, B. a. Keagy, and S. J. Burnham, “Prediction of wrist arteriovenous fistula maturation with preoperative vein mapping with ultrasonography,” *Journal of Vascular Surgery*, vol. 36, no. 3, pp. 460–463, 2002.

- [123] J. van der Linden, T. W. Lameris, A. H. van den Meiracker, A. A. E. A. de Smet, P. J. Blankestijn, and M. A. van den Dorpel, "Forearm venous distensibility predicts successful arteriovenous fistula.," *American journal of kidney diseases : the official journal of the National Kidney Foundation*, vol. 47, pp. 1013–9, jun 2006.
- [124] T. S. Huber, C. K. Ozaki, T. C. Flynn, W. A. Lee, S. a. Berceli, C. M. Hirneise, L. M. Carlton, J. W. Carter, E. a. Ross, and J. M. Seeger, "Prospective validation of an algorithm to maximize native arteriovenous fistulae for chronic hemodialysis access," *Journal of Vascular Surgery*, vol. 36, no. 3, pp. 452–459, 2002.
- [125] W. F. Weitzel, K. Kim, D. W. Park, J. Hamilton, M. Odonnell, T. J. Cichonski, and J. M. Rubin, "High-Resolution Ultrasound Elasticity Imaging to Evaluate Dialysis Fistula Stenosis," in *Seminars in Dialysis*, vol. 22, pp. 84–89, Wiley Online Library, 2009.
- [126] W. F. Weitzel, K. Kim, P. K. Henke, and J. M. Rubin, "High-resolution ultrasound speckle tracking may detect vascular mechanical wall changes in peripheral artery bypass vein grafts," *Annals of Vascular Surgery*, vol. 23, no. 2, pp. 201–206, 2009.
- [127] D. W. Park, G. H. Kruger, J. M. Rubin, J. Hamilton, P. Gottschalk, A. J. Shih, and W. F. Weitzel, "In-vivo blood velocity gradient for renal disease patients," *Ultrasound in Medicine & Biology*, 2012.
- [128] D. W. Park, G. H. Kruger, J. M. Rubin, J. Hamilton, P. Gottschalk, R. E. Dodde, A. J. Shih, and W. F. Weitzel, "Quantification of Ultrasound Correlation-based Flow Velocity Mapping and Edge Velocity Gradient Measurement," pp. 1–49, 2013.
- [129] D. W. Park, G. H. Kruger, J. M. Rubin, J. Hamilton, P. Gottschalk, R. E. Dodde, A. J. Shih, and W. F. Weitzel, "In vivo vascular wall shear rate and circumferential strain of renal disease patients.," *Ultrasound in medicine & biology*, vol. 39, no. 2, pp. 241–52, 2013.
- [130] D. W. Park, M. S. Richards, J. M. Rubin, J. Hamilton, G. H. Kruger, and W. F. Weitzel, "Arterial elasticity imaging: comparison of finite-element analysis models with high-resolution ultrasound speckle tracking," *Cardiovascular Ultrasound*, vol. 8, no. 1, p. 22, 2010.
- [131] C. Tomasi, "Detection and Tracking of Point Features Technical Report CMU-CS-91-132," *Image Rochester NY*, vol. 91, no. April, pp. 1–22, 1991.

- [132] S. P. Wu, S. Ringgaard, S. Oyre, M. S. Hansen, S. Rasmus, and E. M. Pedersen, "Wall Shear Rates Differ between the Normal Carotid, Femoral, and Brachial Arteries: An In Vivo MRI Study," *Journal of Magnetic Resonance Imaging*, vol. 19, no. 2, pp. 188–193, 2004.
- [133] P. V. Stroeve, P. R. Hoskins, and W. J. Easson, "Distribution of wall shear rate throughout the arterial tree: A case study," *Atherosclerosis*, vol. 191, no. 2, pp. 276–280, 2007.
- [134] C. Barry, C. Allott, N. John, P. Mellor, P. Arundel, D. Thomson, and J. Waterton, "Three-dimensional freehand ultrasound: Image reconstruction and volume analysis," *Ultrasound in Medicine & Biology*, vol. 23, no. 8, pp. 1209–1224, 1997.
- [135] R. Rohling, A. Gee, and L. Berman, "Three-dimensional spatial compounding of ultrasound images," *Medical Image Analysis*, vol. 1, no. 3, pp. 177–193, 1997.
- [136] R. W. Prager, a. Gee, and L. Berman, "Stradx: real-time acquisition and visualization of freehand three-dimensional ultrasound," *Medical image analysis*, vol. 3, no. 2, pp. 129–140, 1999.
- [137] R. Rohling, a. Gee, and L. Berman, "A comparison of freehand three-dimensional ultrasound reconstruction techniques," *Medical image analysis*, vol. 3, no. 4, pp. 339–359, 1999.
- [138] O. V. Solberg, F. Lindseth, H. Torp, R. E. Blake, and T. A. Nagelhus Hernes, "Freehand 3D Ultrasound Reconstruction Algorithms A Review," *Ultrasound in Medicine & Biology*, vol. 33, no. 7, pp. 991–1009, 2007.
- [139] K. Konner, "The Arteriovenous Fistula," *Journal of the American Society of Nephrology*, vol. 14, no. 6, pp. 1669–1680, 2003.
- [140] J. Dewey C. F., S. R. Bussolari, J. Gimbrone M. A., and P. F. Davies, "The Dynamic Response of Vascular Endothelial Cells to Fluid Shear Stress," *Journal of Biomechanical Engineering*, vol. 103, no. 3, pp. 177–185, 1981.
- [141] B. J. Ballermann, a. Dardik, E. Eng, and a. Liu, "Shear stress and the endothelium," *Kidney international. Supplement*, vol. 67, pp. S100–S108, 1998.
- [142] K. Jozwik and D. Obidowski, "Numerical simulations of the blood flow through vertebral arteries," *Journal of Biomechanics*, vol. 43, no. 2, pp. 177–185, 2010.

- [143] D. Jodko, D. Obidowski, P. Reorowicz, and K. Jóźwik, “Simulations of the blood flow in the arterio-venous fistula for haemodialysis,” *Acta of Bioengineering and Biomechanics*, vol. 16, no. 1, pp. 69–74, 2014.
- [144] S. Korsheed, M. T. Eldehni, S. G. John, R. J. Fluck, and C. W. McIntyre, “Effects of arteriovenous fistula formation on arterial stiffness and cardiovascular performance and function.,” *Nephrology, dialysis, transplantation : official publication of the European Dialysis and Transplant Association - European Renal Association*, vol. 26, no. 10, pp. 3296–302, 2011.
- [145] T. Khamdaeng, J. Luo, J. Vappou, P. Terdtoon, and E. E. Konofagou, “Arterial stiffness identification of the human carotid artery using the stress-strain relationship in vivo,” *Ultrasonics*, vol. 52, no. 3, pp. 402–411, 2012.
- [146] W. Osika, F. Dangardt, J. Grönros, U. Lundstam, A. Myredal, M. Johansson, R. Volkmann, T. Gustavsson, L.-M. Gan, and P. Friberg, “Increasing peripheral artery intima thickness from childhood to seniority.,” *Arteriosclerosis, thrombosis, and vascular biology*, vol. 27, no. 3, pp. 671–676, 2007.
- [147] Y. O. Kim, Y. J. Choi, J. I. Kim, Y. S. Kim, B. S. Kim, C. W. Park, H. C. Song, S. A. Yoon, Y. S. Chang, and B. K. Bang, “The impact of intima-media thickness of radial artery on early failure of radiocephalic arteriovenous fistula in hemodialysis patients.,” *Journal of Korean medical science*, vol. 21, no. 2, pp. 284–9, 2006.
- [148] A. R. Pries, B. Reglin, and T. W. Secomb, “Remodeling of blood vessels: Responses of diameter and wall thickness to hemodynamic and metabolic stimuli,” *Hypertension*, vol. 46, no. 4, pp. 725–731, 2005.
- [149] S. Laurent, X. Girerd, J. J. Mourad, P. Lacolley, L. Beck, P. Boutouyrie, J. P. Mignot, and M. Safar, “Elastic modulus of the radial artery wall material is not increased in patients with essential hypertension.,” *Arteriosclerosis, Thrombosis, and Vascular Biology*, vol. 14, no. 7, pp. 1223–1231, 1994.
- [150] A. C. Burton, “Relation of structure to function of the tissues of the wall of blood vessels.,” *Physiological reviews*, vol. 34, no. 4, pp. 619–642, 1954.
- [151] G. a. Holzapfel, T. C. Gasser, and R. W. Ogden, “A new constitutive framework for arterial wall mechanics and a comparative study of material models,” *J. Elasticity*, vol. 61, pp. 1–48, 2000.

- [152] D. Seif, T. Mailhot, P. Perera, and D. Mandavia, "Caval Sonography in Shock," *Journal of Ultrasound in Medicine*, vol. 31, no. 12, pp. 1885–1890, 2012.
- [153] T. Kusaba, K. Yamaguchi, and H. Oda, "Echography of the inferior vena cava for estimating fluid removal from patients undergoing hemodialysis," *Japanese Journal of Nephrology*, vol. 38, no. 3, pp. 119–123, 1996.
- [154] C. Barbier, Y. Loubières, C. Schmit, J. Hayon, J. L. Ricôme, F. Jardin, and A. Vieillard-Baron, "Respiratory changes in inferior vena cava diameter are helpful in predicting fluid responsiveness in ventilated septic patients," *Intensive Care Medicine*, vol. 30, pp. 1740–1746, 2004.
- [155] Z. Zhang, X. Xu, S. Ye, and L. Xu, "Ultrasonographic measurement of the respiratory variation in the inferior vena cava diameter is predictive of fluid responsiveness in critically ill patients: systematic review and meta-analysis.," *Ultrasound in medicine & biology*, vol. 40, pp. 845–53, may 2014.
- [156] D. J. Funk, E. Jacobsohn, and A. Kumar, "Role of the venous return in critical illness and shock: part II-shock and mechanical ventilation.," *Critical care medicine*, vol. 41, pp. 573–9, 2013.
- [157] F. Michard and J.-l. Teboul, "Predicting Fluid Responsiveness in ICU Patients * A Critical Analysis of the Evidence," *Critical Care Review*, 2001.
- [158] T. Sonoo, K. Nakamura, T. Ando, K. Sen, A. Maeda, E. Kobayashi, I. Sakuma, K. Doi, S. Nakajima, and N. Yahagi, "Prospective analysis of cardiac collapsibility of inferior vena cava using ultrasonography," *Journal of Critical Care*, pp. 18–21, 2015.
- [159] D. J. Wallace, M. Allison, and M. B. Stone, "Inferior vena cava percentage collapse during respiration is affected by the sampling location: An ultrasound study in healthy volunteers," *Academic Emergency Medicine*, vol. 17, no. 1, pp. 96–99, 2010.
- [160] D. J. Blehar, D. Resop, B. Chin, M. Dayno, and R. Gaspari, "Inferior vena cava displacement during respirophasic ultrasound imaging.," *Critical ultrasound journal*, vol. 4, no. 1, p. 18, 2012.
- [161] P. E. Shrout and J. L. Fleiss, "Intraclass correlations: uses in assessing rater reliability.," 1979.

- [162] K. S. Cole, "Electric Impedance of Suspensions of Spheres," *The Journal of General Physiology*, pp. 29–36, 1928.
- [163] K. S. Cole, "Electric Phase Angle of Cell Membranes," *The Journal of general physiology*, vol. 15, no. 6, pp. 641–649, 1932.
- [164] K. S. Cole, "Electric Impedance of Hipponee Eggs," *The Journal of General Physiology*, vol. 18, no. 6, pp. 877–887, 1935.
- [165] K. S. Cole, "Electric impedance of marine egg membranes," *Transactions of the Faraday Society*, vol. 33, no. 966, p. 966, 1937.
- [166] K. S. Cole and H. J. Curtis, "Electric Impedance of Single Marine Eggs," *The Journal of General Physiology*, pp. 591–599, 1937.
- [167] K. S. Cole and R. H. Cole, "Dispersion and Absorption in Dielectrics: I. Alternating Current Characteristics," *Journal of Chemical Physics*, vol. 9, pp. 341–351, 1941.
- [168] K. S. Cole and R. H. Cole, "Dispersion and Absorption in Dielectrics: II. Direct Current Characteristics," *Journal of Chemical Physics*, vol. 10, pp. 98–105, 1942.
- [169] C. Gabriel, C. Gabriel, S. Gabriel, S. Gabriel, E. Corthout, and E. Corthout, "The dielectric properties of biological tissues: I. Literature survey.," *Physics in medicine and biology*, vol. 41, no. 11, pp. 2231–49, 1996.
- [170] S. Gabriel, R. W. Lau, and C. Gabriel, "The Dielectric Properties of Biological Tissues: II. Measurements in the frequency range 10 Hz to 20 GHz," *Phys. Med. Biol.*, vol. 41, pp. 2251–2269, 1996.
- [171] S. Gabriel, R. W. Lau, and C. Gabriel, "The dielectric properties of biological tissues: III. Parametric models for the dielectric spectrum of tissues.," *Physics in medicine and biology*, vol. 41, no. 11, pp. 2271–2293, 1996.
- [172] G. Martinsen, S. Clausen, J. B. Nysæther, and S. Grimnes, "Utilizing characteristic electrical properties of the epidermal skin layers to detect fake fingers in biometric fingerprint systems - A pilot study," *IEEE Transactions on Biomedical Engineering*, vol. 54, no. 5, pp. 891–894, 2007.
- [173] Y. Ebina, M. Ekida, and H. Hashimoto, "During the Growth and Fermentation," vol. 33, pp. 1290–1295, 1989.

- [174] M. Pérez, R. Muñiz, C. Torre, B. García, C. Carleos, R. Crespo, and L. Cárcel, “Impedance spectrometry for monitoring alcoholic fermentation kinetics under wine-making industrial conditions,” *XIX IMEKO World Congress Fundamental and Applied Metrology*, pp. 2574–2578, 2009.
- [175] J. Lepetit, P. Salé, R. Favier, and R. Dalle, “Electrical impedance and tenderisation in bovine meat.,” *Meat science*, vol. 60, no. 1, pp. 51–62, 2002.
- [176] J. Kerry, J. Kerry, and D. Ledward, *Meat processing: Improving quality (Vol. 76)*. 2002.
- [177] T. Sun, S. Gawad, C. Bernabini, N. G. Green, and H. Morgan, “Broadband single cell impedance spectroscopy using maximum length sequences: theoretical analysis and practical considerations,” *Measurement Science and Technology*, vol. 18, no. 9, pp. 2859–2868, 2007.
- [178] I. K’Owino and O. Sadik, “Impedance Spectroscopy: A Powerful Tool for Rapid Biomolecular Screening and Cell Culture Monitoring,” *Electroanalysis*, vol. 17, no. 23, pp. 2101–2113, 2005.
- [179] D. Miklavcic, N. Pavselj, and F. X. Hart, “Electric properties of tissues,” *Wiley encyclopedia of biomedical engineering*, vol. 209, no. 922834401, pp. 1–12, 2006.
- [180] K. S. Paulson, M. K. Pidcock, and C. N. McLeod, “A probe for organ impedance measurement.,” *IEEE transactions on bio-medical engineering*, vol. 51, no. 10, pp. 1838–1844, 2004.
- [181] K. R. Foster and H. C. Lukaski, “Whole-body impedancewhat does it measure?,” *Clinical Nutrition*, vol. 64, no. 3 Suppl, pp. 388S–396S, 1996.
- [182] R. F. Tanabe, Z. M. A. de Azevedo, V. M. Fonseca, M. V. M. Peixoto, L. A. dos Anjos, M. I. C. Gaspar-Elsas, D. C. B. C. Moore, and E. G. Ramos, “Distribution of bioelectrical impedance vector values in multi-ethnic infants and pre-school children,” *Clinical Nutrition*, vol. 31, no. 1, pp. 144–148, 2012.
- [183] B. Belmont, R. E. Dodde, and A. J. Shih, “Impedance of tissue-mimicking phantom material under compression,” *Journal of Electrical Bioimpedance*, vol. 4, pp. 2–12, feb 2013.
- [184] M. Y. Jaffrin and H. Morel, “Body fluid volumes measurements by impedance: A review of bioimpedance spectroscopy (BIS) and bioimpedance analysis (BIA) methods,” *Medical Engineering and Physics*, vol. 30, pp. 1257–1269, 2008.

- [185] G. M. Chertow, J. M. Lazarus, N. L. Lew, L. Ma, and E. G. Lowrie, "Bioimpedance norms for the hemodialysis population.," *Kidney international*, vol. 52, no. 6, pp. 1617–1621, 1997.
- [186] U. G. Kyle, I. Bosaeus, A. D. De Lorenzo, P. Deurenberg, M. Elia, J. M. Gómez, B. L. Heitmann, L. Kent-Smith, J. C. Melchior, M. Pirlich, H. Scharfetter, A. M. W. J. Schols, and C. Pichard, "Bioelectrical impedance analysis - Part I: Review of principles and methods," *Clinical Nutrition*, vol. 23, no. 5, pp. 1226–1243, 2004.
- [187] U. G. Kyle, I. Bosaeus, A. D. De Lorenzo, P. Deurenberg, M. Elia, J. Manuel Gómez, B. Lilienthal Heitmann, L. Kent-Smith, J.-C. Melchior, M. Pirlich, H. Scharfetter, A. M.W.J Schols, and C. Pichard, "Bioelectrical impedance analysis part II: utilization in clinical practice," *Clinical Nutrition*, vol. 23, no. 6, pp. 1430–1453, 2004.
- [188] a. Lozano, J. Rosell, and R. Pallás-Areny, "A multifrequency multichannel electrical impedance data acquisition system for body fluid shift monitoring.," *Physiological measurement*, vol. 16, no. 4, pp. 227–37, 1995.
- [189] a. Piccoli, S. Nigrelli, a. Caberlotto, S. Bottazzo, B. Rossi, L. Pillon, and Q. Maggiore, "Bivariate normal values of the bioelectrical impedance vector in adult and elderly populations.," *American Journal of Clinical Nutrition*, vol. 61, no. 2, pp. 269–270, 1995.
- [190] K. Cha, G. M. Chertow, J. Gonzalez, J. M. Lazarus, and D. W. Wilmore, "Multifrequency bioelectrical impedance estimates the distribution of body water.," *Journal of applied physiology (Bethesda, Md. : 1985)*, vol. 79, no. 4, pp. 1316–1319, 1995.
- [191] H. C. Lukaski, "Biological indexes considered in the derivation of the bioelectrical impedance analysis.," *The American journal of clinical nutrition*, vol. 64, no. 3 Suppl, pp. 397S–404S, 1996.
- [192] A. Thomasset, "Measurement of the extracellular fluid volume by the electrochemical method. Biophysical significance of 1 kilocycle impedance of the human body," *Lyon medical*, vol. 214, pp. 131–143, aug 1965.
- [193] J. Nyboer, "Electrical impedance plethysmography; a physical and physiologic approach to peripheral vascular study.," *Circulation*, vol. 2, pp. 811–821, 1950.
- [194] W. G. Kubicek, J. N. Karnegis, R. P. Patterson, D. A. Witsoe, and R. H. Mattson, "Development and evaluation of an impedance cardiac output system," *Aerospace medicine*, vol. 37, pp. 1208–1212, dec 1966.

- [195] W. G. Kubicek, "On the source of peak first time derivative (dZ/dt) during impedance cardiography.," *Annals of biomedical engineering*, vol. 17, no. 5, pp. 459–462, 1989.
- [196] a. Sherwood, M. T. Allen, J. Fahrenberg, R. M. Kelsey, W. R. Lovallo, and L. J. P. Van Doornen, "Methodological guidelines for impedance cardiography," 1990.
- [197] I. Mohammed and C. Phillips, "Techniques for Determining Cardiac Output in the Intensive Care Unit," *Critical Care Clinics*, vol. 26, no. 2, pp. 355–364, 2010.
- [198] H. Keren, D. Burkhoff, and P. Squara, "Evaluation of a noninvasive continuous cardiac output monitoring system based on thoracic bioreactance.," *American journal of physiology. Heart and circulatory physiology*, vol. 293, no. 1, pp. H583–H589, 2007.
- [199] S. Grimnes and Ø. G. Martinsen, *Bioimpedance and Bioelectricity Basics*. Academic Press, 3 ed., 2014.
- [200] M. Cheney, D. Isaacson, and J. C. Newel, "Electrical impedance tomography," *SIAM RReview*, vol. 41, no. 1, pp. 85–101, 1999.
- [201] V. Bekos and J. J. Marini, "Monitoring the Mechanically Ventilated Patient," *Critical Care Clinics*, vol. 23, no. 3, pp. 575–611, 2007.
- [202] C. Putensen, H. Wrigge, and J. Zinserling, "Electrical impedance tomography guided ventilation therapy.," *Current opinion in critical care*, vol. 13, no. 3, pp. 344–350, 2007.
- [203] T. Sato, M. Nishinaga, a. Kawamoto, T. Ozawa, and H. Takatsuji, "Accuracy of a continuous blood pressure monitor based on arterial tonometry.," *Hypertension*, vol. 21, no. 6 Pt 1, pp. 866–874, 1993.
- [204] M. T. Clay and T. C. Ferree, "Weighted regularization in electrical impedance tomography with applications to acute cerebral stroke.," *IEEE transactions on medical imaging*, vol. 21, no. 6, pp. 629–37, 2002.
- [205] C. H. Xu, L. Wang, X. T. Shi, F. S. You, F. Fu, R. G. Liu, M. Dai, Z. W. Zhao, G. D. Gao, and X. Z. Dong, "Real-time imaging and detection of intracranial haemorrhage by electrical impedance tomography in a piglet model.," *The Journal of International Medical Research*, vol. 38, no. 5, pp. 1596–604, 2010.

- [206] D. B. Geselowitz, "An application of electrocardiographic lead theory to impedance plethysmography.," *IEEE transactions on bio-medical engineering*, vol. 18, no. 1, pp. 38–41, 1971.
- [207] J. Lehr, "A Vector Derivation Useful in Impedance Plethysmographic Field Calculations," *IEEE Transactions on Biomedical Engineering*, no. March, pp. 156–157, 1972.
- [208] P. Pelosi, M. Gama de Abreu, and P. R. M. Rocco, "New and conventional strategies for lung recruitment in acute respiratory distress syndrome.," *Critical care (London, England)*, vol. 14, no. 2, p. 210, 2010.
- [209] K. M. Mahmoud and A. S. Ammar, "A comparison between two different alveolar recruitment maneuvers in patients with acute respiratory distress syndrome," *International Journal of Critical Illness and Injury Science*, vol. 1, no. 2, pp. 114–120, 2011.
- [210] M. Badet, F. Bayle, J.-C. Richard, and C. Guérin, "Comparison of optimal positive end-expiratory pressure and recruitment maneuvers during lung-protective mechanical ventilation in patients with acute lung injury/acute respiratory distress syndrome.," *Respiratory care*, vol. 54, no. 7, pp. 847–854, 2009.
- [211] X. Monnet, D. Osman, C. Ridet, B. Lamia, C. Richard, and J.-L. Teboul, "Predicting volume responsiveness by using the end-expiratory occlusion in mechanically ventilated intensive care unit patients.," *Critical care medicine*, vol. 37, pp. 951–956, mar 2009.
- [212] X. Monnet and J. L. Teboul, "Assessment of volume responsiveness during mechanical ventilation: recent advances," *Critical Care*, vol. c, no. Chapter 32, pp. 385–396, 2013.
- [213] T. Dietel, G. Filler, R. Grenda, and N. Wolfish, "Bioimpedance and inferior vena cava diameter for assessment of dialysis dry weight," *Pediatric Nephrology*, vol. 14, pp. 903–907, 2000.
- [214] J. M. Brennan, "Handcarried Ultrasound Measurement of the Inferior Vena Cava for Assessment of Intravascular Volume Status in the Outpatient Hemodialysis Clinic," *Clinical Journal of the American Society of Nephrology*, vol. 1, no. 4, pp. 749–753, 2006.

- [215] M. J. Osypka and D. P. Bernstein, “Electrophysiologic principles and theory of stroke volume determination by thoracic electrical bioimpedance.,” 1999.
- [216] L. M. Tierney, M. a. Whooley, and S. Saint, “Oxygen saturation: a fifth vital sign?,” *The Western journal of medicine*, vol. 166, no. 4, pp. 285–286, 1997.
- [217] M. K. Merboth and S. Barnason, “Managing pain: the fifth vital sign,” *The Nursing clinics of North America*, vol. 35, pp. 375–383, jun 2000.
- [218] T. L. Petty, “The forgotten vital signs.,” apr 1994.
- [219] D. Evans, B. Hodgkinson, and J. Berry, “Vital signs in hospital patients: A systematic review,” *International Journal of Nursing Studies*, vol. 38, no. 6, pp. 643–650, 2001.
- [220] Geriatrics and Extended Care Strategic Healthcare Group National Pain Management Coordinating Committee, “Pain as the 5th vital sign toolkit,” *Pain*, no. October, 2000.
- [221] R. A. Mularski, F. White-Chu, D. Overbay, L. Miller, S. M. Asch, and L. Ganzini, “Measuring pain as the 5th vital sign does not improve quality of pain management,” *Journal of General Internal Medicine*, vol. 21, no. 6, pp. 607–612, 2006.
- [222] K. a. Lorenz, C. D. Sherbourne, L. R. Shugarman, L. V. Rubenstein, L. Wen, A. Cohen, J. R. Goebel, E. Hagenmeier, B. Simon, A. Lanto, and S. M. Asch, “How reliable is pain as the fifth vital sign?,” *Journal of the American Board of Family Medicine : JABFM*, vol. 22, no. 3, pp. 291–298, 2009.
- [223] S. C. Millasseau, J. M. Ritter, K. Takazawa, and P. J. Chowienczyk, “Contour analysis of the photoplethysmographic pulse measured at the finger.,” *Journal of hypertension*, vol. 24, no. 8, pp. 1449–1456, 2006.
- [224] X. Teng and Y. Zhang, “Continuous and noninvasive estimation of arterial blood pressure using a photoplethysmographic approach,” *Proceedings of the 25th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (IEEE Cat. No.03CH37439)*, vol. 4, pp. 3153–3156, 2003.
- [225] a. L. Pauca, M. F. O’Rourke, and N. D. Kon, “Prospective evaluation of a method for estimating ascending aortic pressure from the radial artery pressure waveform.,” *Hypertension*, vol. 38, no. 4, pp. 932–937, 2001.

- [226] F. Seoane, J. Ferreira, J. J. Sánchez, and R. Bragós, “An analog front-end enables electrical impedance spectroscopy system on-chip for biomedical applications.,” *Physiological measurement*, vol. 29, pp. S267–78, jun 2008.
- [227] J. Ferreira, F. Seoane, a. Ansele, and R. Bragos, “AD5933-based spectrometer for electrical bioimpedance applications,” *Journal of Physics: Conference Series*, vol. 224, p. 012011, 2010.
- [228] U. Pliquet and A. Barthel, “Interfacing the AD5933 for bio-impedance measurements with front ends providing galvanostatic or potentiostatic excitation,” *Journal of Physics: Conference Series*, vol. 407, p. 012019, 2012.
- [229] C. Margo, J. Katrib, M. Nadi, and a. Rouane, “A four-electrode low frequency impedance spectroscopy measurement system using the AD5933 measurement chip.,” *Physiological measurement*, vol. 34, no. 4, pp. 391–405, 2013.
- [230] Analog Devices, “AD5933: 1 MSPS , 12-Bit Impedance Converter , Network Analyzer,” *Datasheet*, p. 40, 2013.
- [231] Texas Instruments, “Low-Cost , Integrated Analog Front-End for Weight-Scale and Body Composition Measurement AFE4300,” no. June 2012, 2013.
- [232] H. P. Schwan, “Linear and nonlinear electrode polarization and biological materials,” *Annals of Biomedical Engineering*, vol. 20, no. 3, pp. 269–288, 1992.
- [233] H. P. Schwan, “Electrode polarization impedance and measurements in biological materials,” *Annals New York Academy Sciences*, vol. 148, no. A1, pp. 191–209, 1968.
- [234] S. Rossi, M. Pessione, V. Radicioni, G. Baglione, M. Vatteroni, P. Dario, and L. Della Torre, “A low power bioimpedance module for wearable systems,” *Sensors and Actuators A: Physical*, vol. 232, pp. 359–367, 2015.
- [235] M. E. S. Luna, S. Primas, J. M. Stivoric, C. Singleton, P. Savalia, P. Panchalan, S. Nabanja, S. H.-Y. Cheng, I. Mahammad, and S. Sharma, “Physiological information generation based on bioimpedance signals,” 2015.